



Release Notes:

Version F.05.17 Operating System

for the HP Procurve Switches 2512 and 2524

These release notes include information on the following:

- Downloading switch software and Documentation from the Web (Page 1)
- Enhancements in Release F.05.xx (Page 5)
- Enhancements in Release F.04.08 (Page 57)
- Enhancements in Release F.02.11 (Page 137)
- Enhancements in Release F.02.02 (Page 153)
- Updates and corrections for the *Management and Configuration Guide* (page 223)
- Software fixes for Series 2500 switch software releases (page 229)

Caution: Archive Pre-F.05.17 Configuration Files

A configuration file saved while using release F.05.17 or later software is not backward-compatible with earlier software versions. For this reason, HP recommends that you archive the most recent configuration on switches using software releases earlier than F.05.17 before you update any switches to software release F.05.17 or later.

© Copyright 2001, 2003 Hewlett-Packard Company, LP.

This document contains information which is protected by copyright. Reproduction, adaptation, or translation without prior permission is prohibited, except as allowed under the copyright laws.

Publication Number

5990-3102
June 2003
Edition 2

Applicable Product

HP ProCurve Switch 2512 (J4812A)
HP ProCurve Switch 2524 (J4813A)

Trademark Credits

Microsoft, Windows, Windows 95, and Microsoft Windows NT are registered trademarks of Microsoft Corporation.

Software Credits

SSH in the HP Procurve Series 2500 switches is based on the OpenSSH software toolkit. For more information on OpenSSH, visit <http://www.openssh.com>.

Disclaimer

The information contained in this document is subject to change without notice.

HEWLETT-PACKARD COMPANY MAKES NO WARRANTY OF ANY KIND WITH REGARD TO THIS MATERIAL, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. Hewlett-Packard shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this material.

Hewlett-Packard Company shall not be liable for technical or editorial errors or omissions contained herein. The information is provided "as is" without warranty of any kind and is subject to change without notice. The warranties for Hewlett-Packard Company products are set forth in the express limited warranty statements for such products. Nothing herein should be construed as constituting an additional warranty.

Hewlett-Packard assumes no responsibility for the use or reliability of its software on equipment that is not furnished by Hewlett-Packard.

Warranty

See the Customer Support/Warranty booklet included with the product.

A copy of the specific warranty terms applicable to your Hewlett-Packard products and replacement parts can be obtained from your HP Sales and Service Office or authorized dealer.

Hewlett-Packard Company
8000 Foothills Boulevard, m/s 5552
Roseville, California 95747-5552
<http://www.hp.com/go/hpprocurve>

Contents

Software Management	1
Downloading Switch Documentation and Software	1
Downloading Software to the Switch	2
TFTP Download from a Server	2
Xmodem Download From a PC or Unix Workstation	3
Saving Configurations While Using the CLI	4
HP ProCurve Switch Software Key	4
Enhancements in Release F.05.xx	5
Clarification of Time Zone Issue	5
Isolated Port Groups (Enhanced)	6
Options for Isolated Port Groups	6
Operating Rules for Port Isolation	8
Configuring Port Isolation on the Switch	8
Steps for Configuring Port Isolation	8
Configuring and Viewing Port-Isolation	9
Messages Related to Port-Isolation Operation	13
Troubleshooting Port-Isolation Operation	13
Configuring Port-Based Access Control (802.1x)	14
Overview	14
Why Use Port-Based Access Control?	14
General Features	14
How 802.1x Operates	16
Authenticator Operation	16
Switch-Port Supplicant Operation	16
Terminology	18
General Operating Rules and Notes	19
General Setup Procedure for Port-Based Access Control (802.1x)	20
Do These Steps Before You Configure 802.1x Operation	20
Overview: Configuring 802.1x Authentication on the Switch	21
Configuring Switch Ports as 802.1x Authenticators	22

1. Enable 802.1x Authentication on Selected Ports	23
3. Configure the 802.1x Authentication Method	26
4. Enter the RADIUS Host IP Address(es)	27
5. Enable 802.1x Authentication on the Switch	27
802.1x Open VLAN Mode	28
Introduction	28
Use Models for 802.1x Open VLAN Modes	29
Operating Rules for Authorized-Client and Unauthorized-Client VLANs	32
Setting Up and Configuring 802.1x Open VLAN Mode	34
802.1x Open VLAN Operating Notes	38
Option For Authenticator Ports: Configure Port-Security To Allow Only 802.1x Devices	39
Configuring Switch Ports To Operate As Supplicants for 802.1x Connections to Other Switches	41
Displaying 802.1x Configuration, Statistics, and Counters	45
Show Commands for Port-Access Authenticator	45
Viewing 802.1x Open VLAN Mode Status	47
Show Commands for Port-Access Supplicant	50
How RADIUS/802.1x Authentication Affects VLAN Operation	51
Messages Related to 802.1x Operation	54
IGMP Version 3 Support	55
Enhancements in Release F.04.08	57
Using Friendly (Optional) Port Names	58
Configuring and Operating Rules for Friendly Port Names	58
Configuring Friendly Port Names	59
Displaying Friendly Port Names with Other Port Data	60
Configuring Secure Shell (SSH)	63
Terminology	65
Prerequisite for Using SSH	65
Public Key Format Requirement	65
Steps for Configuring and Using SSH for Switch and Client Authentication	66
General Operating Rules and Notes	68
Configuring the Switch for SSH Operation	69
Further Information on SSH Client Public-Key Authentication	80
Messages Related to SSH Operation	85

Troubleshooting SSH Operation	86
Configuring RADIUS Authentication and Accounting	89
Terminology	90
Switch Operating Rules for RADIUS	90
General RADIUS Setup Procedure	91
Configuring the Switch for RADIUS Authentication	92
Configuring RADIUS Accounting	102
Operating Rules for RADIUS Accounting	103
Viewing RADIUS Statistics	108
Changing RADIUS-Server Access Order	113
Messages Related to RADIUS Operation	114
Troubleshooting RADIUS Operation	115
IP Preserve: Retaining VLAN-1 IP Addressing Across Configuration File Downloads	117
Configuring Port-Based Priority for Incoming Packets	120
Messages Related to Prioritization	123
Troubleshooting Prioritization	123
Using the "Kill" Command To Terminate Remote Sessions	124
Configuring Rapid Reconfiguration Spanning Tree (RSTP)	125
Overview	125
Transitioning from STP to RSTP	126
Configuring RSTP	127
Enhancements in Release F.02.11	137
Fast-Uplink Spanning Tree Protocol (STP)	137
The Show Tech Command for Listing Switch Configuration and Operating Details	151
Enhancements in Release F.02.02	153
Documentation for Enhancements in Release F.02.02	153
TACACS+ Authentication for Centralized Control of Switch Access Security	154
Series 2500 Switch Authentication Options	155
Terminology Used in TACACS Applications:	156
General System Requirements	157
TACACS+ Operation	158
General Authentication Setup Procedure	158

Configuring TACACS+ on the Switch	161
Viewing the Switch's Current Authentication Configuration	162
Viewing the Switch's Current TACACS+ Server Contact Configuration	162
Configuring the Switch's Authentication Methods	163
Configuring the Switch's TACACS+ Server Access	166
How Authentication Operates	170
General Authentication Process Using a TACACS+ Server	170
Local Authentication Process	171
Using the Encryption Key	172
General Operation	172
Encryption Options in the Switch	172
Controlling Web Browser Interface Access When Using TACACS+ Authentication	173
Messages	174
Operating Notes	174
Troubleshooting TACACS+ Operation	175
CDP	177
Introduction	177
CDP Terminology	178
General CDP Operation	178
Outgoing Packets	179
Incoming CDP Packets	179
Configuring CDP on the Switch	182
Viewing the Switch's Current CDP Configuration	182
Viewing the Current Contents of the Switch's CDP Neighbors Table	183
Clearing (Resetting) the CDP Neighbors Table	184
Configuring CDP Operation	185
Effect of Spanning Tree (STP) On CDP Packet Transmission	187
How CDP Selects the CDP Neighbor's IP Address When Multiple VLANs Are Present	187
CDP Neighbor Data and MIB Objects	188
CDP Operating Notes	190
Troubleshooting CDP Operation	190
New Time Synchronization Protocol Options	192

TimeP Time Synchronization	192
SNTP Time Synchronization	193
Overview: Selecting a Time Synchronization Protocol or Turning Off Time Protocol Operation	193
General Steps for Running a Time Protocol on the Switch:	193
Disabling Time Synchronization	194
SNTP: Viewing, Selecting, and Configuring	194
Menu: Viewing and Configuring SNTP	195
CLI: Viewing and Configuring SNTP	197
TimeP: Viewing, Selecting, and Configuring	202
Menu: Viewing and Configuring TimeP	203
CLI: Viewing and Configuring TimeP	205
SNTP Unicast Time Polling with Multiple SNTP Servers	209
Address Prioritization	209
Adding and Deleting SNTP Server Addresses	210
Menu Interface Operation with Multiple SNTP Server Addresses Configured	211
SNTP Messages in the Event Log	211
Operation and Enhancements for Multimedia Traffic Control (IGMP)	212
How Data-Driven IGMP Operates	212
IGMP Operates With or Without IP Addressing	213
Fast-Leave IGMP	214
Forced Fast-Leave IGMP	216
Configuration Options for Forced Fast-Leave	216
CLI: Listing the Forced Fast-Leave Configuration	216
CLI: Configuring Per-Port Forced Fast-Leave IGMP	218
Querier Operation	218
The Switch Excludes Well-Known or Reserved Multicast Addresses from IP Multicast Filtering	219
Switch Memory Operation	220
Port Security: Changes to Retaining Learned Static Addresses Across a Reboot	220
Recommended Port Security Procedures	220
Retention of Static Addresses	221
Username Assignment and Prompt	222

Updates and Corrections for the Management and Configuration Guide	223
Changes in Commands for Viewing the Current Configuration Files	223
Change in CLI Command for Listing Intrusion Alerts	224
Changes for Listing Port and Trunk Group Statistics	224
Time Protocol Changes	224
Change in Command Line (CLI) Operation	224
Restoring the Factory-Default Configuration, Including Usernames and Passwords ...	225
Incomplete IP Multicast (IGMP) Filtering Data	225
GVRP Does Not Require a Common VLAN	225
Incomplete Information on Saving Configuration Changes	226
Update to Information on Duplicate MAC Addresses Across VLANs	226
Incorrect Command Listing for Viewing Configuration Files	227
New and Corrected Information on Primary VLAN Usage	227
Misleading Statement About VLANs	228
 Software Fixes	 229
 Series 2500 FAQs from the HP ProCurve Website	 248
Auto-Negotiation	248
Cabling	249
General	249
Features	250
Spanning Tree Protocol	251
VLANs	253
Troubleshooting	255

Software Management

Caution: Archive Pre-F.05.17 Configuration Files

A configuration file saved while using release F.05.17 or later software is not backward-compatible with earlier software versions. For this reason, HP recommends that you archive the most recent configuration on switches using software releases earlier than F.05.17 before you update any switches to software release F.05.17 or later.

Downloading Switch Documentation and Software

You can download software version F.05.17 and the corresponding product documentation from HP's ProCurve website as described below.

To Download a Software Version:

1. Go to HP's ProCurve website at <http://www.hp.com/go/hpprocurve>.
2. Click on **software** (in the sidebar).
3. Under **latest software**, click on **switches**.

To Download Product Documentation: You will need the Adobe® Acrobat® Reader to view, print, and/or copy the product documentation.

1. Go to HP's ProCurve website at <http://www.hp.com/go/hpprocurve>.
2. Click on **technical support**, then **manuals**.
3. Click on the name of the product for which you want documentation.
4. On the resulting web page, double-click on a document you want.
5. When the document file opens, click on the disk icon  in the Acrobat® toolbar and save a copy of the file.

Downloading Software to the Switch

HP periodically provides switch operating system (OS) updates through the HP ProCurve website (<http://www.hp.com/go/hpprocurve>). After you acquire the new OS file, you can use one of the following methods for downloading the operating system (OS) code to the switch:

- For a TFTP transfer from a server, do either of the following:
 - Click on **Download OS** in the Main Menu of the switch's menu interface and use the (default) **TFTP** option.
 - Use the `copy tftp` command in the switch's CLI (see below).
- For an Xmodem transfer from a PC or Unix workstation, do either of the following:
 - Click on **Download OS** in the Main Menu of the switch's menu interface and select the **Xmodem** option.
 - Use the `copy xmodem` command in the switch's CLI (page 3).
- The software update utility included in some network management applications.
- A switch-to-switch file transfer

Note

Downloading a new OS does not change the current switch configuration. The switch configuration is contained in a separate file that can also be transferred, for example, for archive purposes or to be used in another switch of the same model. See also .

This section describes how to use the CLI to download an OS to the switch. You can also use the menu interface for OS downloads. For more information, refer to the *Management and Configuration Guide* for the Series 2500 switches.

TFTP Download from a Server

Syntax: `copy tftp flash <ip-address> <remote-os-file>`

For example, to download an OS file named `F_05_17.swi` from a TFTP server with the IP address of 10.28.227.103:

1. Execute the copy command as shown below:

```
HP2512# copy tftp flash 10.28.227.103 F_05_17.swi
Device will be rebooted, do you want to continue [y/n]? y
00224K _
```

2. When the switch finishes downloading the OS file from the server, it displays this progress message:

Validating and Writing System Software to FLASH . . .

3. After the switch reboots, it displays the CLI or Main Menu, depending on the **Logon Default** setting last configured in the menu's Switch Setup screen.

Xmodem Download From a PC or Unix Workstation

This procedure assumes that:

- The switch is connected via the Console RS-232 port on a PC operating as a terminal. (Refer to the Installation Guide you received with the switch for information on connecting a PC as a terminal and running the switch console interface.)
- The switch operating system (OS) 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 Windows NT terminal emulator, you would use the **Send File** option in the **Transfer** dropdown menu.)

Syntax: `copy xmodem flash <unix | pc>`

For example, to download an OS file named F_05.17.swi from a PC:

1. To reduce the download time, you may want to increase the baud rate in your terminal emulator and in the switch to a value such as 57600 bits per second. (The baud rate must be the same in both devices.) For example, to change the baud rate in the switch to 57600, execute this command:

```
HP2512(config)# console baud-rate 57600
```

(If you use this option, be sure to set your terminal emulator to the same baud rate.)

2. Execute the following command in the CLI:

```
HP2512(config)# copy xmodem flash pc
Device will be rebooted, do you want to continue [y/n]? y
Press 'Enter' and start XMODEM on your host...
```

3. Execute the terminal emulator commands to begin the Xmodem transfer.

The download can take several minutes, depending on the baud rate used in the transfer.

When the download finishes, the switch automatically reboots itself and begins running the new OS version.

4. To confirm that the operating system downloaded correctly:

```
HP2512> show system
```

Check the **Firmware revision** line.

Software Management

Saving Configurations While Using the CLI

5. If you increased the baud rate on the switch (step 1), use the same command to return it to its previous setting. (HP recommends a baud rate of 9600 bits per second for most applications.)
(Remember to return your terminal emulator to the same baud rate as the switch.)

Saving Configurations While Using the CLI

The switch operates with two configuration files:

- **Running-Config File:** Exists in volatile memory and controls switch operation. Rebooting the switch erases the current running-config file and replaces it with an exact copy of the current startup-config file. To save a configuration change, you must save the running configuration to the startup-config file.
- **Startup-Config File:** Exists in flash (non-volatile) memory and preserves the most recently-saved configuration as the "permanent" configuration. When the switch reboots for any reason, an exact copy of the current startup-config file becomes the new running-config file in volatile memory.

When you use the CLI to make a configuration change, the switch places the change in the running-config file. If you want to preserve the change across reboots, you must save the change to the startup-config file. Otherwise, the next time the switch reboots, the change will be lost. There are two ways to save configuration changes while using the CLI:

- Execute the **write memory** command from the Manager, Global, or Context configuration level.
- When exiting from the CLI to the Main Menu, press **[Y]** (for Yes) when you see the **save configuration** prompt:

```
Do you want to save current configuration [y/n] ?
```

HP ProCurve Switch Software Key

Letter	HP ProCurve Switch
C	1600M, 2400M, 2424M, 4000M, 8000M
E	Series 5300XL (5304XL and 5308XL)
F	Series 2500 (2512 and 2524)
G	Series 4100GL (4104GL and 4108GL)
H	Switch 2650 and Switch 6108

Enhancements in Release F.05.xx

Enhancement	Summary	Page
Isolated Port Groups	Originally added in release F.04.08 to provide an alternative to VLANs, this feature now offers two new isolation groups: group1 and group2.	6
Port-Based Access Control (802.1x) with Open VLAN Mode	Originally added in release F.04.08 to provide access control through a RADIUS server, this feature now includes Open VLAN Mode. This gives you a means for allowing a client computer without 802.1x supplicant software to temporarily join an unauthorized-client VLAN and proceed with initialization services, such as acquiring IP addressing, 802.1x supplicant software, and other optional services you may want to provide.	14
IGMP Version 3 Support	The switch now supports operation with IGMPv3 traffic.	55

Clarification of Time Zone Issue

Starting with the F.05.xx version of the switch operating system software, the method of configuring the Time Zone for TimeP or SNTP configuration has been updated. Previous switch software, for all HP ProCurve switches, used positive time offset values for time zones that are West of GMT and negative values for time zones that are East of GMT. The standards indicate that time zones West of GMT should be designated by negative offset values, and time zones East of GMT by positive values. Software version F.05.xx updates this configuration method, but if you use the same values for indicating time zones as you did for previous HP ProCurve switches, the time will be set incorrectly on your HP ProCurve Switches 2512 and 2524. For example, for previous HP ProCurve switches, the US Pacific time zone was configured by entering **+480**. With software version F.05.xx, the US Pacific time zone must now be configured by entering **-480**.

Isolated Port Groups (Enhanced)

Isolated Port-Group Commands

[no] port-isolation	page 9
port-isolation [ethernet] < port-list > mode < uplink public group1 group2 private local >	page 9
show port-isolation	page 9

The Isolated Port Groups feature originally included in release F.04.08 has been enhanced in release F.05.xx with the inclusion of two new port isolation groups (**group1** and **group2**).

Isolated port groups provide an alternative to VLANs for isolating end nodes on your network, while simplifying network administration. This feature enables you to isolate traffic to and from specific end-node devices, which enhances security and also helps in such areas as selectively preventing internet use. There are, however, some limitations, as outlined in the "Rules of Operation", described later in this section.

Caution

The Isolated Port Groups feature is intended for rare situations where using VLANs is not possible. This feature can interfere with other switch features, and improper configuration will result in unexpected connectivity problems. Refer to “Operating Rules for Port Isolation” on page 8.

The Isolated Port Groups feature operates within the context of the individual switch. It does not restrict free communication on the designated uplink port(s) to other devices on the network. A node connected to any type of port (group1, group2, private, etc.) on one Series 2500 switch can communicate with a node connected to any type of port (group1, group2, private, etc.) on another Series 2500 switch if the two switches are connected through their uplink ports.

Options for Isolated Port Groups

Using Isolated Port Groups, you can control traffic between ports on the switch by assigning an appropriate port type to each port. The options include:

- Uplink (the default)
 - Public
 - Group1
 - Group2
 - Private
 - Local
-

When you configure isolated port groups on a switch, traffic is allowed to move between the switch ports as described in table 1 and shown in figure 1, both below.

Table 1. Communication Allowed Between Port-Isolation Types within a Switch

Port Type:	Permits Traffic To and From This Port Type?						Notes
	Uplink Ports	Public Ports	Group1 Ports	Group2 Ports	Local Ports	Private Ports	
Public Ports	Yes	Yes	No	No	Yes	No	Typical switch ports: For intra-switch operation, allows communication among end nodes on public and local ports, and between end nodes on public ports and the uplink port(s).
Uplink Ports	Yes	Yes	Yes	Yes	No	Yes	Allows communication between uplink ports and end nodes on public and private ports. Uplink ports are intended for connecting the switch to the network core. When you enable port isolation on the switch, Uplink is the default port-isolation mode setting for individual ports.
Group1 Ports	Yes	No	Yes	No	No	No	Allows communication among end nodes on other group-1 ports, and between end nodes on Group1 ports and the Uplink port(s).
Group2 Ports	Yes	No	No	Yes	No	No	Allows communication among end nodes on other Group2 ports, and between end nodes on Group2 ports and the Uplink port(s).
Local Ports	No	Yes	No	No	Yes	No	Allows communication among end nodes on local and public ports.
Private Ports	Yes	No	No	No	No	No	Allows communication only between end nodes and uplink ports.

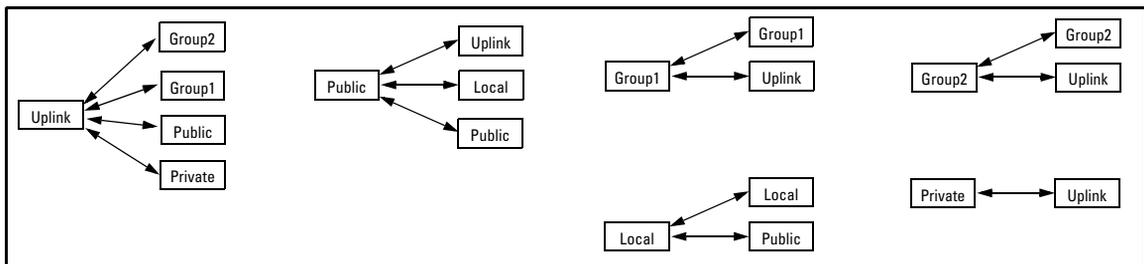


Figure 1. Communication Allowed Between Port-Isolation Types within a Switch

Operating Rules for Port Isolation

- Port Isolation is intended only for networks that do not use VLAN tagging. (The switch must be in the default VLAN configuration before you configure port-isolation.)
- Multiple VLANs are not allowed on the switch. If multiple VLANs exist on the switch, delete them and return the ports to the original default configuration as untagged members of VLAN 1. (VLAN configuration changes are not supported if port-isolation is running on the switch.)
- Trunking is supported *only* on Uplink ports between switches. Remove any other port trunking from the switch.
- LACP is allowed only on the Uplink ports. *For security, LACP (active or passive) must be disabled on all other ports on the switch.* To disable LACP active or passive on the switch's ports, use this command syntax:

```
no int e <port-numbers> lacp
```
- GVRP must be disabled (the default).
- IGMP operates only in non-data-driven mode, and works only on uplink ports. The switch floods multicast IP traffic arriving at non-uplink ports.
- A Series 2500 switch with port-isolation enabled cannot export its port-isolation configuration. However, a Series 2500 configuration file on a server can include port-isolation commands.
- The Isolated Port Groups feature operates within the context of the individual switch. It does not restrict free communication on the designated uplink port(s) to other devices on the network. A node connected to any non-local port (group1, group2, private, etc.) on one Series 2500 switch can communicate with a node connected to any non-local port (group1, group2, private, etc.) on another Series 2500 switch if the two switches are connected through their uplink ports.
- Enabling port isolation and configuring individual ports to specific, non-default modes are separate steps. You must first enable port isolation. When you do so, all ports are configured in the (default) **Uplink** mode.

Configuring Port Isolation on the Switch

Steps for Configuring Port Isolation

1. Remove all non-default VLANs from the switch and ensure that all ports are untagged members of the default VLAN (VID = 1).
2. Identify the devices you will connect to the switch's ports.

3. Configure all equipment you plan to attach to the switch (such as servers and other switches) to eliminate VLAN tagging on ports connected to the Series 2500 switch(es) on which you are using Port Isolation.
4. Determine the mode assignment you want for each port on the switch. (When you enable port-isolation, the switch configures all ports to the default **Uplink** mode.)
5. Remove port trunks you have configured from ports that you plan to configure in public, local, or private mode.
6. Disable LACP on all ports that you plan to configure in public, local, or private mode. To do so, use this command: **no interface e < port-list > lacp**.
7. Enable port isolation on the switch.
8. Configure the non-default port-isolation mode for each port that you do not want to operate in the **Uplink** mode.
9. Connect the switch ports to the other devices in your port-isolation plan.
10. Test the operation of all ports you are using for links to the other devices.
11. When you are satisfied that your port-isolation configuration is working properly, execute **write mem** to store the configuration in the startup-config file.

Configuring and Viewing Port-Isolation

Syntax: [no] port-isolation

*Without any port-list or mode parameters, enables port isolation on the switch and sets all ports to the Uplink mode. The **no** version disables port isolation and also causes all individual ports to be set to the (default) Uplink mode the next time you enable port isolation.*

[ethernet] < port-list > mode < uplink | public | group1 | group2 | private | local >

*Specifies the ports you want to configure to a particular port-isolation mode (**uplink**—the default—**public, group1, group2, private, local**).*

show port-isolation

Lists the switch's port-isolation status and, if enabled, the port-isolation mode and which ports, if any, are in a port trunk.

show running-config

Lists the switch's running configuration, including port-isolation settings.

show config

Lists the switch's startup configuration, including port-isolation settings.

Note

The **no port-isolation** command erases all port-isolation mode settings from memory. This means that whenever you disable, then re-enable port isolation, all ports on the switch will be set to the (default) **Uplink** mode.

For example, suppose that the switch is in its default configuration (no multiple VLANs; GVRP disabled, all ports untagged members of the default VLAN—VID = 1) with two optional gigabit transceivers installed, and you wanted to use the switch ports as shown in table table 2, “Port Isolation Plan”:

Table 2. Port Isolation Plan

Port	Use	Allowed	Traffic Blocked
1 - 3 1	Local ports only for isolated workgroup access. (No network or internet access.)	<ul style="list-style-type: none"> Traffic between any ports in the local set (ports 1, 2, and 3) Traffic between any port in the local set and any port in the public set (ports 10, 11, or 12) 	Traffic between any port in the local set and any port in the private, group1, or uplink port sets
4 - 8 2	Group1 ports for workgroup and network/internet access	<ul style="list-style-type: none"> Traffic between any ports in the group1 set (ports 4 through 8) Traffic between any port in group1 and the uplink ports 	Traffic between any port in the group1 set (ports 4 - 8) and any public, private, or local ports
9 3	Private port to a secure end node; no traffic exchange with non-uplink ports on the switch.	Traffic between port 9 (private) and the gigabit trunk used as an uplink (ports 13 and 14).	Traffic between port 9 and any port in the local, public, or group1 port sets, or any other private port on the switch
10 - 12 4	Public ports for typical end-node access.	<ul style="list-style-type: none"> Traffic between any ports in the public set (ports 10, 11, and 12) Traffic between any port in the public set and any port in the local or uplink port sets 	Traffic between any port in the public set (ports 10 - 12) and any port in the group1 or private port sets
13 - 14 5	Gigabit uplink to the network.	<ul style="list-style-type: none"> Traffic between any ports in the uplink set (ports 13 and 14) Traffic between any port in the uplink set and any port in the public, private, or group1 sets 	Traffic between any port in the uplink set and any port in the local set

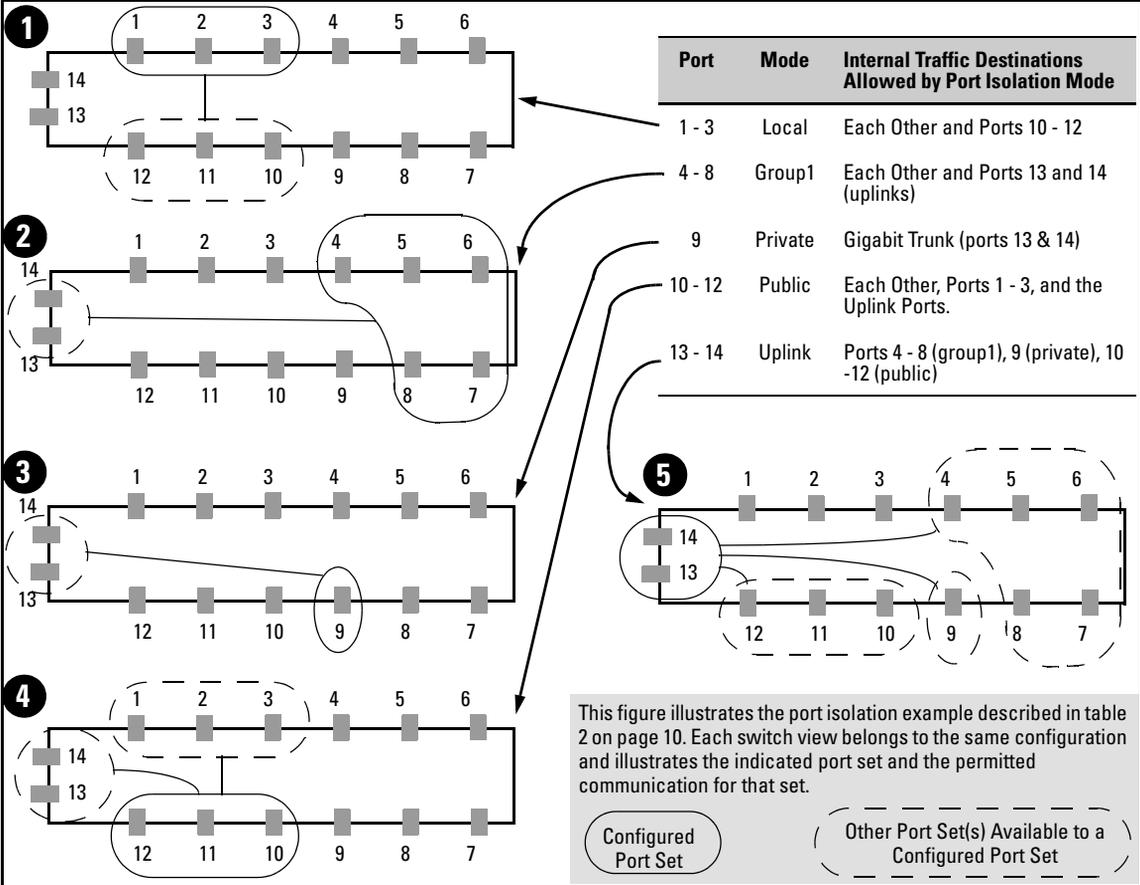


Figure 2. Example of Isolating Ports on a Series 2500 Switch

Assuming a switch in the factory-default configuration, you would configure the port isolation plan in figure 2 as follows:

```
HP2512(config)# no interface e 1-12 lacp
HP2512(config)# port-isolation
```

****** CAUTION ******

Contact your local Customer Care Center before activating this feature to receive proper configuration instructions. Failure to configure this feature properly will result in unexpected connectivity problems.

```
Continue with enabling port isolation? y
HP2512(config)# port-isolation e 1-3 mode local
HP2512(config)# port-isolation e 4-8 mode group1
HP2512(config)# port-isolation e 9 mode private
HP2512(config)# port-isolation e 10-12 mode public

HP2512(config)# show port-isolation
Port Isolation: Enabled

Port Isolation

Port  Mode
-----
1     Local
2     Local
3     Local
4     Group1
5     Group1
6     Group1
7     Group1
8     Group1
9     Private
10    Public
11    Public
12    Public
13    Uplink
14    Uplink
```

Remember to disable LACP on ports that will be configured for Public, Group1, Group2, Private, or Local mode. (Refer to "Operating Rules for Port Isolation" on page 8.)

When you enter the command to enable port isolation, the switch displays a caution and prompts you to indicate how to proceed. Type **[Y]** to continue with enabling port isolation; **[N]** to leave port isolation disabled. See the Caution on page 6.

Uplink mode is the default setting for all ports when you enable port-isolation. Since these two ports were not explicitly configured, above, they remain in the Uplink mode (and do not need to be explicitly configured as uplinks).

Figure 3. Example of Port-Isolation Configuration

Messages Related to Port-Isolation Operation

Message	Meaning
Port Isolation is disabled. It must be enabled first.	In the switch's factory-default state or after you execute no port-isolation , you must enable port isolation (by executing port-isolation alone) before entering commands for changing the mode on one or more ports.

Troubleshooting Port-Isolation Operation

Symptom	Possible Cause
Connectivity problems.	<ul style="list-style-type: none"> • A port may be configured as a tagged member of a VLAN, or multiple VLANs may be configured on the switch. Ensure that all ports are untagged members of VLAN 1 (the default VLAN) and that no other VLANs are configured on the switch. • Illegal port trunking. Port Isolation does not allow trunks on Private ports, or more than one Port-Isolation type in a trunk. Also, Port Isolation allows an LACP trunk only on Uplink ports. • A port on a device connected to the switch may be configured as a tagged member of a VLAN. • GVRP may be enabled on the switch. <p>See "Operating Rules for Port Isolation" on page 8 and "Steps for Configuring Port Isolation" on page 8.</p>

Configuring Port-Based Access Control (802.1x)

Overview

Feature	Default	Menu	CLI	Web
Configuring Switch Ports as 802.1x Authenticators	Disabled	n/a	page -22	n/a
Configuring 802.1x Open VLAN Mode	Disabled	n/a	page -28	n/a
Configuring Switch Ports to Operate as 802.1x Supplicants	Disabled	n/a	page -41	n/a
Displaying 802.1x Configuration, Statistics, and Counters	n/a	n/a	page -45	n/a
How 802.1x Affects VLAN Operation	n/a	n/a	page -51	n/a
RADIUS Authentication and Accounting	Refer to "Configuring RADIUS Authentication and Accounting" on page 89			

Why Use Port-Based Access Control?

Local Area Networks are often deployed in a way that allows unauthorized clients to attach to network devices, or allows unauthorized users to get access to unattended clients on a network. Also, the use of DHCP services and zero configuration make access to networking services easily available. This exposes the network to unauthorized use and malicious attacks. While access to the network should be made easy, uncontrolled and unauthorized access is usually not desirable. 802.1x provides access control along with the ability to control user profiles from a central RADIUS server while allowing users access from multiple points within the network.

General Features

802.1x on the Series 2500 switches includes the following:

- Switch operation as both an authenticator (for supplicants having a point-to-point connection to the switch) and as a supplicant for point-to-point connections to other 802.1x-aware switches.
 - Authentication of 802.1x clients using a RADIUS server and either the EAP or CHAP protocol.
 - Provision for enabling clients that do not have 802.1 supplicant software to use the switch as a path for downloading the software and initiating the authentication process (802.1x Open VLAN mode).
 - Supplicant implementation using CHAP authentication and independent username and password configuration on each port.
- Prevention of traffic flow in either direction on unauthorized ports.
- Local authentication of 802.1x clients using the switch's local username and password (as an alternative to RADIUS authentication).

- Temporary on-demand change of a port's VLAN membership status to support a current client's session. (This does not include ports that are members of a trunk.)
- Session accounting with a RADIUS server, including the accounting update interval.
- Use of Show commands to display session counters.
- With port-security enabled for port-access control, limit a port to one 802.1x client session at a given time.

Authenticating Users. Port-Based Access Control (802.1x) provides switch-level security that allows LAN access only to users who enter the authorized RADIUS username and password on 802.1x-capable clients (supplicants). This simplifies security management by allowing you to control access from a master database in a single server (although you can use up to three RADIUS servers to provide backups in case access to the primary server fails). It also means a user can enter the same username and password pair for authentication, regardless of which switch is the access point into the LAN. Note that you can also configure 802.1x for authentication through the switch's local username and password instead of a RADIUS server, but doing so increases the administrative burden, decentralizes username/password administration, and reduces security by limiting authentication to one Operator/Manager password set for all users.

Providing a Path for Downloading 802.1x Supplicant Software. For clients that do not have the necessary 802.1x supplicant software, there is also the option to configure the 802.1x Open VLAN mode. This mode allows you to assign such clients to an isolated VLAN through which you can provide the necessary supplicant software these clients need to begin the authentication process. (Refer to "802.1x Open VLAN Mode" on page 28.)

Authenticating One Switch to Another. 802.1x authentication also enables the switch to operate as a supplicant when connected to a port on another switch running 802.1x authentication.

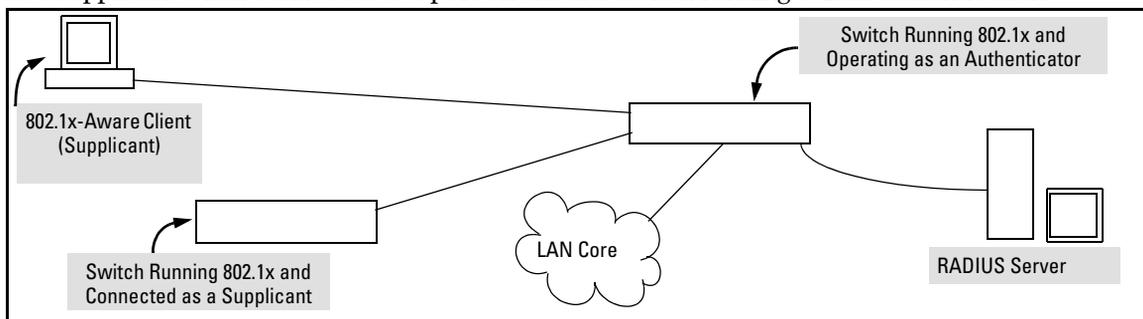


Figure 4. Example of an 802.1x Application

Accounting . The Series 2500 switches also provide RADIUS Network accounting for 802.1x access. Refer to "Configuring RADIUS Authentication and Accounting" on page 89.

How 802.1x Operates

Authenticator Operation

This operation provides security on a direct, point-to-point link between a single client and the switch, where both devices are 802.1x-aware. (If you expect desirable clients that do not have the necessary 802.1x supplicant software, you can provide a path for downloading such software by using the 802.1x Open VLAN mode—refer to “802.1x Open VLAN Mode” on page 28.) For example, suppose that you have configured a port on the switch for 802.1x authentication operation. If you then connect an 802.1x-aware client (supplicant) to the port and attempt to log on:

1. When the switch detects the client on the port, it blocks access to the LAN from that port.
2. The switch responds with an identity request.
3. The client responds with a user name that uniquely defines this request for the client.
4. The switch responds in one of the following ways:
 - If 802.1x (port-access) on the switch is configured for RADIUS authentication, the switch then forwards the request to a RADIUS server.
 - i. The server responds with an access challenge which the switch forwards to the client.
 - ii. The client then provides identifying credentials (such as a user certificate), which the switch forwards to the RADIUS server.
 - iii. The RADIUS server then checks the credentials provided by the client.
 - iv. If the client is successfully authenticated and authorized to connect to the network, then the server notifies the switch to allow access to the client. Otherwise, access is denied and the port remains blocked.
 - If 802.1x (port-access) on the switch is configured for local authentication, then:
 - i. The switch compares the client’s credentials with the username and password configured in the switch (Operator or Manager level).
 - ii. If the client is successfully authenticated and authorized to connect to the network, then the switch allows access to the client. Otherwise, access is denied and the port remains blocked.

Switch-Port Supplicant Operation

This operation provides security on links between 802.1x-aware switches. For example, suppose that you want to connect two switches, where:

- Switch “A” has port 1 configured for 802.1x supplicant operation.
- You want to connect port 1 on switch “A” to port 5 on switch “B”.

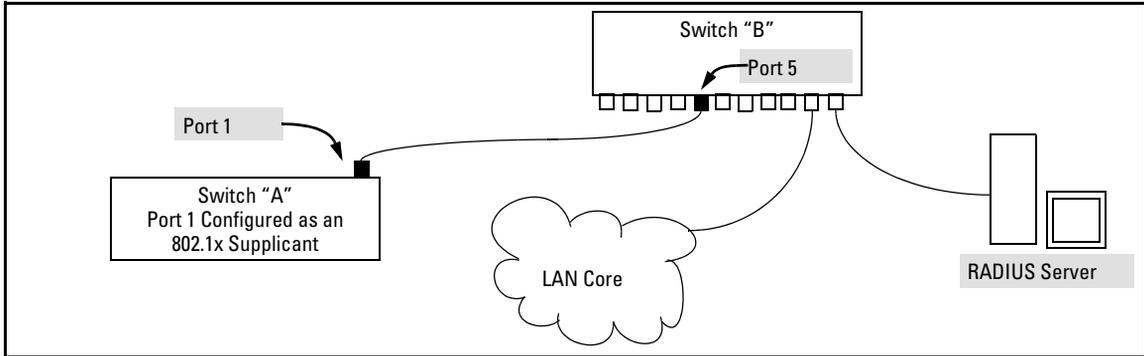


Figure 5. Example of Supplicant Operation

1. When port 1 on switch "A" is first connected to a port on switch "B", or if the ports are already connected and either switch reboots, port 1 begins sending start packets to port 5 on switch "B".
 - If, after the supplicant port sends the configured number of start packets, it does not receive a response, it assumes that switch "B" is not 802.1x-aware, and transitions to the authenticated state. If switch "B" is operating properly and is not 802.1x-aware, then the link should begin functioning normally, but without 802.1x security.
 - If, after sending one or more start packets, port 1 receives a request packet from port 5, then switch "B" is operating as an 802.1x authenticator. The supplicant port then sends a response/ID packet. Switch "B" forwards this request to a RADIUS server.
2. The RADIUS server then responds with an MD5 access challenge that switch "B" forwards to port 1 on switch "A".
3. Port 1 replies with an MD5 hash response based on its username and password or other unique credentials. Switch "B" forwards this response to the RADIUS server.
4. The RADIUS server then analyzes the response and sends either a "success" or "failure" packet back through switch "B" to port 1.
 - A "success" response unblocks port 5 to normal traffic from port 1.
 - A "failure" response continues the block on port 5 and causes port 1 to wait for the "held-time" period before trying again to achieve authentication through port 5.

Note

You can configure a switch port to operate as both a supplicant and an authenticator at the same time.

Terminology

802.1x-Aware: Refers to a device that is running either 802.1x authenticator software or 802.1x client software and is capable of interacting with other devices on the basis of the IEEE 802.1x standard.

Authorized-Client VLAN: Like the Unauthorized-Client VLAN, this is a conventional, static VLAN previously configured on the switch by the System Administrator. The intent in using this VLAN is to provide authenticated clients with network services that are not available on either the port's statically configured VLAN memberships or any VLAN memberships that may be assigned during the RADIUS authentication process. While an 802.1x port is a member of this VLAN, the port is untagged. When the client connection terminates, the port drops its membership in this VLAN.

Authentication Server: The entity providing an authentication service to the switch when the switch is configured to operate as an authenticator. In the case of a Series 2500 switch running 802.1x, this is a RADIUS server (unless local authentication is used, in which case the switch performs this function using its own username and password for authenticating a supplicant).

Authenticator: In HP ProCurve switch applications, a device such as a Series 2500 switch that requires a supplicant to provide the proper credentials (username and password) before being allowed access to the network.

CHAP (MD5): Challenge Handshake Authentication Protocol.

Client: In this application, an end-node device such as a management station, workstation, or mobile PC linked to the switch through a point-to-point LAN link.

EAP (Extensible Authentication Protocol): EAP enables network access that supports multiple authentication methods.

EAPOL: Extensible Authentication Protocol Over LAN, as defined in the 802.1x standard.

Friendly Client: A client that does not pose a security risk if given access to the switch and your network.

MD5: An algorithm for calculating a unique digital signature over a stream of bytes. It is used by CHAP to perform authentication without revealing the shared secret (password).

PVID (Port VID): This is the VLAN ID for the untagged VLAN to which an 802.1x port belongs.

Static VLAN: A VLAN that has been configured as “permanent” on the switch by using the CLI **vlan <vid>** command or the Menu interface.

Supplicant: The entity that must provide the proper credentials to the switch before receiving access to the network. This is usually an end-user workstation, but it can be a switch, router, or another device seeking network services.

Tagged VLAN Membership: This type of VLAN membership allows a port to be a member of multiple VLANs simultaneously. If a client connected to the port has an operating system that supports 802.1q VLAN tagging, then the client can access VLANs for which the port is a tagged member. If the client does not support VLAN tagging, then it can access only a VLAN for which the port is an untagged member. (A port can be an untagged member of only one VLAN at a time.) 802.1x Open VLAN mode does not affect a port's tagged VLAN access unless the port is statically configured as a member of a VLAN that is also configured as the Unauthorized-Client or Authorized-Client VLAN. See also “**Untagged VLAN Membership**”.

Unauthorized-Client VLAN: A conventional, static VLAN previously configured on the switch by the System Administrator. It is used to provide access to a client prior to authentication. It should be set up to allow an unauthenticated client to access only the initialization services necessary to establish an authenticated connection, plus any other desirable services whose use by an unauthenticated client poses no security threat to your network. (Note that an unauthenticated client has access to all network resources that have membership in the VLAN you designate as the Unauthorized-Client VLAN.) A port configured to use a given Unauthorized-Client VLAN does not have to be statically configured as a member of that VLAN as long as at least one other port on the switch is statically configured as a tagged or untagged member of the same Unauthorized-Client VLAN.

Untagged VLAN Membership: A port can be an untagged member of only one VLAN. (In the factory-default configuration, all ports on the switch are untagged members of the default VLAN.) An untagged VLAN membership is *required* for a client that does not support 802.1q VLAN tagging. A port can simultaneously have one untagged VLAN membership and multiple tagged VLAN memberships. Depending on how you configure 802.1x Open VLAN mode for a port, a statically configured, untagged VLAN membership may become unavailable while there is a client session on the port. See also “**Tagged VLAN Membership**”.

General Operating Rules and Notes

- When a port on the switch is configured as either an authenticator or supplicant and is connected to another device, rebooting the switch causes a re-authentication of the link.
- When a port on the switch is configured as an authenticator, it will block access to a client that either does not provide the proper authentication credentials or is not 802.1x-aware. (You can use the optional 802.1x Open VLAN mode to open a path for downloading 802.1x supplicant software to a client, which enables the client to initiate the authentication procedure. Refer to “802.1x Open VLAN Mode” on page 28.)
- If a port on switch “A” is configured as an 802.1x supplicant and is connected to a port on another switch, “B”, that is not 802.1x-aware, access to switch “B” will occur without 802.1x security protection.
- You can configure a port as both an 802.1x authenticator *and* an 802.1x supplicant.

- If a port on switch “A” is configured as both an 802.1x authenticator *and* supplicant and is connected to a port on another switch, “B”, that is not 802.1x-aware, access to switch “B” will occur without 802.1x security protection, but switch “B” will not be allowed access to switch “A”. This means that traffic on this link between the two switches will flow from “A” to “B”, but not the reverse.
- If a client already has access to a switch port when you configure the port for 802.1x authenticator operation, the port will block the client from further network access until it can be authenticated.
- On a port configured for 802.1x with RADIUS authentication, if the RADIUS server specifies a VLAN for the supplicant and the port is a trunk member, the port will be blocked. If the port is later removed from the trunk, the port will try to authenticate the supplicant. If authentication is successful, the port becomes unblocked. Similarly, if the supplicant is authenticated and later the port becomes a trunk member, the port will be blocked. If the port is then removed from the trunk, it tries to re-authenticate the supplicant. If successful, the port becomes unblocked.
- To help maintain security, 802.1x and LACP cannot both be enabled on the same port. If you try to configure 802.1x on a port already configured for LACP (or the reverse) you will see a message similar to the following:

Error configuring port X: LACP and 802.1x cannot be run together.

Note on 802.1x and LACP

To help maintain security, the switch does not allow 802.1x and LACP to both be enabled at the same time on the same port. Refer to “802.1x Operating Messages” on page 54.

General Setup Procedure for Port-Based Access Control (802.1x)

Do These Steps Before You Configure 802.1x Operation

1. Configure a local username and password on the switch for both the Operator (login) and Manager (enable) access levels. (While this may or may not be required for your 802.1x configuration, HP recommends that you use a local username and password pair at least until your other security measures are in place.)
2. Determine which ports on the switch you want to operate as authenticators and/or supplicants, and disable LACP on these ports. (See the “Note on 802.1x and LACP” on page 20.)

3. Determine whether to use the optional 802.1x Open VLAN mode for clients that are not 802.1x-aware; that is, for clients that are not running 802.1x supplicant software. (This will require you to provide downloadable software that the client can use to enable an authentication session.) For more on this topic, refer to “802.1x Open VLAN Mode” on page 28.
4. For each port you want to operate as a supplicant, determine a username and password pair. You can either use the same pair for each port or use unique pairs for individual ports or subgroups of ports. (This can also be the same local username/password pair that you assign to the switch.)
5. Unless you are using only the switch’s local username and password for 802.1x authentication, configure at least one RADIUS server to authenticate access requests coming through the ports on the switch from external supplicants (including switch ports operating as 802.1x supplicants). You can use up to three RADIUS servers for authentication; one primary and two backups. Refer to the documentation provided with your RADIUS application.

Overview: Configuring 802.1x Authentication on the Switch

This section outlines the steps for configuring 802.1x on the switch. For detailed information on each step, refer to “Configuring RADIUS Authentication and Accounting” on page 89 or “Configuring Switch Ports To Operate As Supplicants for 802.1x Connections to Other Switches” on page 41.

1. Enable 802.1x authentication on the individual ports you want to serve as authenticators. On the ports you will use as authenticators, either accept the default 802.1x settings or change them, as necessary. Note that, by default, the port-control parameter is set to **auto** for all ports on the switch. This requires a client to support 802.1x authentication and to provide valid credentials to get network access. Refer to page -23.
2. If you want to provide a path for clients without 802.1x supplicant software to download the software so that they can initiate an authentication session, enable the 802.1x Open VLAN mode on the ports you want to support this feature. Refer to page 28.
3. Configure the 802.1x authentication type. Options include:
 - Local Operator username and password (the default). This option allows a client to use the switch’s local username and password as valid 802.1x credentials for network access.
 - EAP RADIUS: This option requires your RADIUS server application to support EAP authentication for 802.1x.
 - CHAP (MD5) RADIUS: This option requires your RADIUS server application to support CHAP (MD5) authentication.

See page -26.

4. If you select either **eap-radius** or **chap-radius** for step 3, use the **radius host** command to configure up to three RADIUS server IP address(es) on the switch. See page -27.
5. Enable 802.1x authentication on the switch. See page 23.

6. Test both the authorized and unauthorized access to your system to ensure that the 802.1x authentication works properly on the ports you have configured for port-access.

Note

If you want to implement the optional port security feature (step 7) on the switch, you should first ensure that the ports you have configured as 802.1x authenticators operate as expected.

7. If you are using Port Security on the switch, configure the switch to allow only 802.1x access on ports configured for 802.1x operation, and (if desired) the action to take if an unauthorized device attempts access through an 802.1x port. See page 39.
8. If you want a port on the switch to operate as a supplicant in a connection with a port operating as an 802.1x authenticator on another device, then configure the supplicant operation. (Refer to “Configuring Switch Ports To Operate As Supplicants for 802.1x Connections to Other Switches” on page 41.)

Configuring Switch Ports as 802.1x Authenticators

802.1x Authentication Commands	Page
[no] aaa port-access authenticator < [ethernet] < <i>port-list</i> >	23
[control quiet-period tx-period supplicant-timeout server-timeout max-requests reauth-period auth-vid unauth-vid initialize reauthenticate clear-statistics]	23
aaa authentication port-access < local eap-radius chap-radius >	26
[no] aaa port-access authenticator active	22
[no] port-security [ethernet] < <i>port-list</i> > learn-mode port-access	39
802.1x Open VLAN Mode Commands	28
802.1x Supplicant Commands	41
802.1x-Related Show Commands	45
RADIUS server configuration	27

1. Enable 802.1x Authentication on Selected Ports

This task configures the individual ports you want to operate as 802.1x authenticators for point-to-point links to 802.1x-aware clients or switches. (Actual 802.1x operation does not commence until you perform step 5 on page 21 to activate 802.1x authentication on the switch.)

Note

When you enable 802.1x authentication on a port, the switch automatically disables LACP on that port. However, if the port is already operating in an LACP trunk, you must remove the port from the trunk before you can configure it for 802.1x authentication.

Syntax: `aaa port-access authenticator < port-list >`

Enables specified ports to operate as 802.1x authenticators with current per-port authenticator configuration. To activate configured 802.1x operation, you must enable 802.1x authentication. Refer to “5. Enable 802.1x Authentication on the switch” on page 21.

[control < authorized | auto | unauthorized >]

Controls authentication mode on the specified port:

authorized: *Also termed Force Authorized. Grants access to any device connected to the port. In this case, the device does not have to provide 802.1x credentials or support 802.1x authentication. (However, you can still configure console, Telnet, or SSH security on the port.)*

auto (the default): *The device connected to the port must support 802.1x authentication and provide valid credentials in order to get network access. (You have the option of using the Open VLAN mode to provide a path for clients without 802.1x supplicant software to download this software and begin the authentication process. Refer to “802.1x Open VLAN Mode” on page 28.)*

unauthorized: *Also termed Force Unauthorized. Do not grant access to the network, regardless of whether the device provides the correct credentials and has 802.1x support. In this state, the port blocks access to any connected device.*

Syntax: aaa port-access authenticator < port-list > *(Syntax Continued)*

[quiet-period < 0 - 65535 >]

*Sets the period during which the port does not try to acquire a supplicant. The period begins after the last attempt authorized by the **max-requests** parameter fails (next page). (Default: 60 seconds)*

[tx-period < 0 - 65535 >]

Sets the period the port waits to retransmit the next EAPOL PDU during an authentication session. (Default: 30 seconds)

[supplicant-timeout < 1 - 300 >]

Sets the period of time the switch waits for a supplicant response to an EAP request. If the supplicant does not respond within the configured time frame, the session times out. (Default: 30 seconds)

[server-timeout < 1 - 300 >]

*Sets the period of time the switch waits for a server response to an authentication request. If there is no response within the configured time frame, the switch assumes that the authentication attempt has timed out. Depending on the current **max-requests** setting, the switch will either send a new request to the server or end the authentication session. (Default: 30 seconds)*

[max-requests < 1 - 10 >]

*Sets the number of authentication attempts that must time-out before authentication fails and the authentication session ends. If you are using the Local authentication option, or are using RADIUS authentication with only one host server, the switch will not start another session until a client tries a new access attempt. If you are using RADIUS authentication with two or three host servers, the switch will open a session with each server, in turn, until authentication occurs or there are no more servers to try. During the **quiet-period** (previous page), if any, you cannot reconfigure this parameter. (Default: 2)*

Syntax: aaa port-access authenticator < port-list > *(Syntax Continued)*

[reauth-period < 1 - 9999999 >]

Sets the period of time after which clients connected must be re-authenticated. When the timeout is set to 0 the reauthentication is disabled (Default: 0 second)

[unauth-vid < vlan-id >]

Configures an existing static VLAN to be the Unauthorized-Client VLAN. This enables you to provide a path for clients without supplicant software to download the software and begin an authentication session. Refer to “802.1x Open VLAN Mode” on page 28.

[auth-vid < vid >]

Configures an existing, static VLAN to be the Authorized-Client VLAN. Refer to “802.1x Open VLAN Mode” on page 28.

[initialize]

*On the specified ports, blocks inbound and outbound traffic and restarts the 802.1x authentication process. This happens only on ports configured with **control auto** and actively operating as 802.1x authenticators. **Note:** If a specified port is configured with **control authorized** and **port-security**, and the port has learned an authorized address, the port will remove this address and learn a new one from the first packet it receives.*

[reauthenticate]

Forces reauthentication (unless the authenticator is in 'HELD' state).

[clear-statistics]

Clears authenticator statistics counters.

3. Configure the 802.1x Authentication Method

This task specifies how the switch will authenticate the credentials provided by a supplicant connected to a switch port configured as an 802.1x authenticator.

Syntax: `aaa authentication port-access < local | eap-radius | chap-radius >`

Determines the type of RADIUS authentication to use.

local Use the switch's local username and password for supplicant authentication.

eap-radius Use EAP-RADIUS authentication. (Refer to the documentation for your RADIUS server.)

chap-radius Use CHAP-RADIUS (MD-5) authentication. (Refer to the documentation for your RADIUS server application.)

For example, to enable the switch to perform 802.1x authentication using one or more EAP-capable RADIUS servers:

```

HPswitch(config)# aaa authentication port-access eap-radius
HPswitch(config)# show auth

```

Status and Counters - Authentication Information

Login Attempts : 3

Access Task	Login Primary	Login Secondary	Enable Primary	Enable Secondary
Console	Local	None	Local	None
Telnet	Local	None	Local	None
Port-Access	EapRadius			
SSH	Local	None	Local	None

Configuration command for EAP-RADIUS authentication.

802.1x (Port-Access) configured for EAP RADIUS authentication.

Figure 6. Example of 802.1x (Port-Access) Authentication

4. Enter the RADIUS Host IP Address(es)

If you selected either **eap-radius** or **chap-radius** for the authentication method, configure the switch to use 1 to 3 RADIUS servers for authentication. The following syntax shows the basic commands. For coverage of all commands related to RADIUS server configuration, refer to “Configuring RADIUS Authentication and Accounting” on page 89.

Syntax: radius host < ip-address >

Adds a server to the RADIUS configuration.

[key < server-specific key-string >]

Optional. Specifies an encryption key for use during authentication (or accounting) sessions with the specified server. This key must match the key used on the RADIUS server. Use this option only if the specified server requires a different key than configured for the global encryption key.

radius-server key < global key-string >

Specifies the global encryption key the switch uses for sessions with servers for which the switch does not have a server-specific key. This key is optional if all RADIUS server addresses configured in the switch include a server-specific encryption key.

5. Enable 802.1x Authentication on the Switch

After configuring 802.1x authentication as described in the preceding four sections, activate it with this command:

Syntax: aaa port-access authenticator active

Activates 802.1x port-access on ports you have configured as authenticators.

802.1x Open VLAN Mode

802.1x Authentication Commands	page 22
802.1x Supplicant Commands	page 42
802.1x Open VLAN Mode Commands	
[no] aaa port-access authenticator [e] < port-list >	page 37
[auth-vid < vlan-id >]	
[unauth-vid < vlan-id >]	
802.1x-Related Show Commands	page 45
RADIUS server configuration	pages 27

This section describes how to use the 802.1x Open VLAN mode to configure unauthorized-client and authorized-client VLANs on ports configured as 802.1x authenticators.

Introduction

Configuring the 802.1x Open VLAN mode on a port changes how the port responds when it detects a new client. In earlier releases, a “friendly” client computer not running 802.1x supplicant software could not be authenticated on a port protected by 802.1x access security. As a result, the port would become blocked and the client could not access the network. This prevented the client from:

- Acquiring IP addressing from a DHCP server
- Downloading the 802.1x supplicant software necessary for an authentication session

The 802.1x Open VLAN mode solves this problem by temporarily suspending the port’s static, untagged VLAN membership and placing the port in a designated *Unauthorized-Client VLAN*. In this state the client can proceed with initialization services, such as acquiring IP addressing and 802.1x software, and starting the authentication process. Following authentication, the port drops its temporary (untagged) membership in the Unauthorized-Client VLAN and joins (or rejoins) *one* of the following as an *untagged* member:

- **1st Priority:** The port joins a VLAN to which it has been assigned by a RADIUS server during authentication.
- **2nd Priority:** If RADIUS authentication does not include assigning a VLAN to the port, then the switch assigns the port to the VLAN entered in the port’s 802.1x configuration as an *Authorized-Client VLAN*, if configured.

- **3rd Priority:** If the port does not have an Authorized-Client VLAN configured, but does have a static, untagged VLAN membership in its configuration, then the switch assigns the port to this VLAN.

If the port is not configured for any of the above, then it must be a tagged member of at least one VLAN. In this case, if the client is capable of operating in a tagged VLAN, then it can access that VLAN. Otherwise, the connection will fail.

Caution

If a port is a tagged member of a statically configured VLAN, 802.1x Open VLAN mode does not prevent unauthenticated client access to such VLANs if the client is capable of operating in a tagged VLAN environment. To avoid possible security breaches, HP recommends that you not allow a tagged VLAN membership on a port configured for 802.1x Open VLAN mode unless you use the tagged VLAN as the Unauthorized-Client VLAN.

Use Models for 802.1x Open VLAN Modes

You can apply the 802.1x Open VLAN mode in more than one way. Depending on your use, you will need to create one or two static VLANs on the switch for *exclusive* use by per-port 802.1x Open VLAN mode authentication:

- **Unauthorized-Client VLAN:** Configure this VLAN when unauthenticated, friendly clients will need access to some services before being authenticated.
- **Authorized-Client VLAN:** Configure this VLAN for authenticated clients when the port is not statically configured as an untagged member of a VLAN you want clients to use, or when the port is statically configured as an untagged member of a VLAN you do not want clients to use. (A port can be configured as untagged on only one VLAN. When an Authorized-Client VLAN is configured, it will always be untagged and will block the port from using a statically configured, untagged membership in another VLAN.)

Table 3. 802.1x Open VLAN Mode Options

802.1x Per-Port Configuration	Port Response
No Open VLAN mode:	The port automatically blocks a client that cannot initiate an authentication session.
Open VLAN mode with both of the following configured:	
Unauthorized-Client VLAN	<ul style="list-style-type: none"> • When the port detects a client, it automatically becomes an untagged member of this VLAN. If you previously configured the port as a static, tagged member of the VLAN, membership temporarily changes to untagged while the client remains unauthenticated. • If the port already has a statically configured, untagged membership in another VLAN, then the port temporarily closes access to this other VLAN while in the Unauthorized-Client VLAN. • To limit security risks, the network services and access available on the Unauthorized-Client VLAN should include only what a client needs to enable an authentication session. If the port is statically configured as a tagged member of any other VLANs, access to these VLANs remains open, even though the client may not be authenticated. Refer to the Caution on page 29.
Authorized-Client VLAN	<ul style="list-style-type: none"> • After the client is authenticated, the port drops membership in the Unauthorized-Client VLAN and becomes an untagged member of this VLAN. <p data-bbox="594 873 1272 986">Note: If RADIUS authentication assigns a VLAN, the port temporarily becomes a member of the RADIUS-assigned VLAN — instead of the Authorized-Client VLAN—while the client is connected.</p> • If the port is statically configured as a tagged member of a VLAN, and this VLAN is used as the Authorized-Client VLAN, then the port temporarily becomes an untagged member of this VLAN when the client becomes authenticated. When the client disconnects, the port returns to tagged membership in this VLAN. • If the port is statically configured as a tagged member of a VLAN that is not used by 802.1x Open VLAN mode, an unauthenticated client capable of operating in tagged VLANs has access to this VLAN. Refer to the Caution on page 29.

802.1x Per-Port Configuration**Port Response**

Open VLAN Mode with Only an Unauthorized-Client VLAN Configured:

- When the port detects a client, it automatically becomes an untagged member of this VLAN. To limit security risks, the network services and access available on this VLAN should include only what a client needs to enable an authentication session. If the port is statically configured as an untagged member of another VLAN, the switch temporarily removes the port from membership in this other VLAN while membership in the Unauthorized-Client VLAN exists.
- After the client is authenticated, and if the port is statically configured as an untagged member of another VLAN, the port's access to this other VLAN is restored.
- If the port is statically configured as a tagged member of a VLAN that is not used by 802.1x Open VLAN mode, an unauthenticated client capable of operating in tagged VLANs can access this VLAN. Refer to the Caution on page 29.

Note: If RADIUS authentication assigns a VLAN to the port, this assignment overrides any statically configured, untagged VLAN membership on the port (while the client is connected).

Open VLAN Mode with Only an Authorized-Client VLAN Configured:

- Port automatically blocks a client that cannot initiate an authentication session.
- If the client successfully completes an authentication session, the port becomes an untagged member of this VLAN.
- If the port is statically configured as a tagged member of any other VLANs, an authenticated client capable of operating in a tagged VLAN environment can access these VLANs.

Note: If RADIUS authentication assigns a VLAN, the port temporarily becomes a member of the RADIUS-assigned VLAN—instead of the Authorized-Client VLAN—while the client is connected.

Operating Rules for Authorized-Client and Unauthorized-Client VLANs

Condition	Rule
Static VLANs used as <i>Authorized-Client</i> or <i>Unauthorized-Client</i> VLANs	These must be configured on the switch before you configure an 802.1x authenticator port to use them. (Use the vlan < vlan-id > command or the VLAN Menu screen in the Menu interface.)
VLAN Assignment Received from a RADIUS Server	If the RADIUS server specifies a VLAN for an authenticated supplicant connected to an 802.1x authenticator port, this VLAN assignment overrides any Authorized-Client VLAN assignment configured on the authenticator port. This is because both VLANs are untagged, and the switch allows only one untagged VLAN membership per-port. For example, suppose you configured port 4 to place authenticated supplicants in VLAN 20. If a RADIUS server authenticates supplicant "A" and assigns this supplicant to VLAN 50, then the port can access VLAN 50 for the duration of the client session. When the client disconnects from the port, then the port drops these assignments and uses only the VLAN memberships for which it is statically configured.
Temporary VLAN Membership During a Client Session	<ul style="list-style-type: none"> • Port membership in a VLAN assigned to operate as the Unauthorized-Client VLAN is temporary, and ends when the client receives authentication or the client disconnects from the port, whichever is first. • Port membership in a VLAN assigned to operate as the Authorized-Client VLAN is also temporary, and ends when the client disconnects from the port. If a VLAN assignment from a RADIUS server is used instead, the same rule applies.
Effect of Unauthorized-Client VLAN session on untagged port VLAN membership	<ul style="list-style-type: none"> • When an unauthenticated client connects to a port that is already configured with a static, untagged VLAN, the switch temporarily moves the port to the Unauthorized-Client VLAN (also untagged). (While the Unauthorized-Client VLAN is in use, the port does not access the static, untagged VLAN.) • When the client either becomes authenticated or disconnects, the port leaves the Unauthorized-Client VLAN and reacquires its untagged membership in the statically configured VLAN.
Effect of Authorized-Client VLAN session on untagged port VLAN membership.	<ul style="list-style-type: none"> • When a client becomes authenticated on a port that is already configured with a static, untagged VLAN, the switch temporarily moves the port to the Authorized-Client VLAN (also untagged). While the Authorized-Client VLAN is in use, the port does not have access to the statically configured, untagged VLAN. • When the authenticated client disconnects, the switch removes the port from the Authorized-Client VLAN and moves it back to the untagged membership in the statically configured VLAN.

Condition	Rule
Multiple Authenticator Ports Using the Same Unauthorized-Client and Authorized-Client VLANs	<p>You can use the same static VLAN as the Unauthorized-Client VLAN for all 802.1x authenticator ports configured on the switch. Similarly, you can use the same static VLAN as the Authorized-Client VLAN for all 802.1x authenticator ports configured on the switch.</p> <p><i>Caution: Do not use the same static VLAN for both the unauthorized and the Authorized-Client VLAN. Using one VLAN for both creates a security risk by defeating the isolation of unauthenticated clients.</i></p>
Effect of Failed Client Authentication Attempt	<p>When there is an Unauthorized-Client VLAN configured on an 802.1x authenticator port, an unauthorized client connected to the port has access only to the network resources belonging to the Unauthorized-Client VLAN. (There can be an exception to this rule if the port is also a tagged member of a statically configured VLAN. Refer to the Caution on page 29.) This access continues until the client disconnects from the port. (If there is no Unauthorized-Client VLAN configured on the authenticator port, the port simply blocks access for any unauthorized client that cannot be authenticated.)</p>
Sources for an IP Address Configuration for a Client Connected to a Port Configured for 802.x Open VLAN Mode	<p>A client can either acquire an IP address from a DHCP server or have a preconfigured, manual IP address before connecting to the switch.</p>
802.1x Supplicant Software for a Client Connected to a Port Configured for 802.1x Open VLAN Mode	<p>A friendly client, without 802.1x supplicant software, connecting to an authenticator port must be able to download this software from the Unauthorized-Client VLAN before authentication can begin.</p>

Note:

If you use the same VLAN as the Unauthorized-Client VLAN for all authenticator ports, unauthenticated clients on different ports can communicate with each other.

Setting Up and Configuring 802.1x Open VLAN Mode

Preparation. This section assumes use of both the Unauthorized-Client and Authorized-Client VLANs. Refer to Table 3 on page 30 for other options.

Before you configure the 802.1x Open VLAN mode on a port:

- Statically configure an “Unauthorized-Client VLAN” in the switch. The only ports that should belong to this VLAN are ports offering services and access you want available to unauthenticated clients. (802.1x authenticator ports do not have to be members of this VLAN.)

Caution

Do not allow any port memberships or network services on this VLAN that would pose a security risk if exposed to an unauthorized client.

- Statically configure an Authorized-Client VLAN in the switch. The only ports that should belong to this VLAN are ports offering services and access you want available to authenticated clients. 802.1x authenticator ports do not have to be members of this VLAN.

Note that if an 802.1x authenticator port is an untagged member of another VLAN, the port’s access to that other VLAN will be temporarily removed while an authenticated client is connected to the port. For example, if:

- i. Port 5 is an untagged member of VLAN 1 (the default VLAN).
- ii. You configure port 5 as an 802.1x authenticator port.
- iii. You configure port 5 to use an Authorized-Client VLAN.

Then, if a client connects to port 5 and is authenticated, port 5 becomes an untagged member of the Authorized-Client VLAN and is temporarily suspended from membership in the default VLAN.

- If you expect friendly clients to connect without having 802.1x supplicant software running, provide a server on the Unauthorized-Client VLAN for downloading 802.1x supplicant software to the client, and a procedure by which the client initiates the download.
- A client must either have a valid IP address configured before connecting to the switch, or download one through the Unauthorized-Client VLAN from a DHCP server. In the latter case, you will need to provide DHCP services on the Unauthorized-Client VLAN.
- Ensure that the switch is connected to a RADIUS server configured to support authentication requests from clients using ports configured as 802.1x authenticators. (The RADIUS server should not be on the Unauthorized-Client VLAN.)

Note that as an alternative, you can configure the switch to use local password authentication instead of RADIUS authentication. However, this is less desirable because it means that all clients use the same passwords and have the same access privileges. Also, you must use 802.1x supplicant software that supports the use of local switch passwords.

Caution

Ensure that you do not introduce a security risk by allowing Unauthorized-Client VLAN access to network services or resources that could be compromised by an unauthorized client.

Configuring General 802.1x Operation: These steps enable 802.1x authentication, and must be done before configuring 802.1x VLAN operation.

1. Enable 802.1x authentication on the individual ports you want to serve as authenticators. (The switch automatically disables LACP on the ports on which you enable 802.1x.) On the ports you will use as authenticators with VLAN Operation, ensure that the (default) port-control parameter is set to **auto**. This setting requires a client to support 802.1x authentication (with 802.1x supplicant operation) and to provide valid credentials to get network access.

Syntax: `aaa port-access authenticator e <port-list> control auto`

Activates 802.1x port-access on ports you have configured as authenticators.

2. Configure the 802.1x authentication type. Options include:

Syntax: `aaa authentication port-access <local | eap-radius | chap-radius >`

Determines the type of RADIUS authentication to use.

local: *Use the switch's local username and password for supplicant authentication (the default).*

eap-radius *Use EAP-RADIUS authentication. (Refer to the documentation for your RADIUS server.*

chap-radius *Use CHAP-RADIUS (MD5) authentication. (Refer to the documentation for your RADIUS server software.)*

3. If you selected either **eap-radius** or **chap-radius** for step 2, use the **radius host** command to configure up to three RADIUS server IP address(es) on the switch.

Syntax: radius host < ip-address >

Adds a server to the RADIUS configuration.

[key < server-specific key-string >]

Optional. Specifies an encryption key for use with the specified server. This key must match the key used on the RADIUS server. Use this option only if the specified server requires a different key than configured for the global encryption key.

radius-server key < global key-string >

Specifies the global encryption key the switch uses for sessions with servers for which the switch does not have a server-specific key. This key is optional if all RADIUS server addresses configured in the switch include a server-specific encryption key.

4. Activate authentication on the switch.

Syntax: aaa port-access authenticator active

Activates 802.1x port-access on ports you have configured as authenticators.

5. Test both the authorized and unauthorized access to your system to ensure that the 802.1x authentication works properly on the ports you have configured for port-access.

Note

If you want to implement the optional port security feature on the switch, you should first ensure that the ports you have configured as 802.1x authenticators operate as expected. Then refer to “Option For Authenticator Ports: Configure Port-Security To Allow Only 802.1x Devices” on page 39.

After you complete steps 1 and 2, the configured ports are enabled for 802.1x authentication (without VLAN operation), and you are ready to configure VLAN Operation.

Configuring 802.1x Open VLAN Mode. Use these commands to actually configure Open VLAN mode. For a listing of the steps needed to prepare the switch for using Open VLAN mode, refer to “Preparation” on page 34.

Syntax: `aaa port-access authenticator [e] <port-list >`
`[auth-vid <vlan-id >]`
 Configures an existing, static VLAN to be the Authorized-Client VLAN.
`[< unauth-vid <vlan-id >]`
 Configures an existing, static VLAN to be the Unauthorized-Client VLAN.

For example, suppose you want to configure 802.1x port-access with Open VLAN mode on ports 10 - 20 and:

- These two static VLANs already exist on the switch:
 - Unauthorized, VID = 80
 - Authorized, VID = 81
- Your RADIUS server has an IP address of 10.28.127.101. The server uses **rad4all** as a server-specific key string. The server is connected to a port on the Default VLAN.
- The switch's default VLAN is already configured with an IP address of 10.28.127.100 and a network mask of 255.255.255.0

```
HPswitch(config)# aaa authentication port-access eap-radius
```

Configures the switch for 802.1x authentication using an EAP-RADIUS server.

```
HPswitch(config)# aaa port-access authenticator 10-20
```

Configures ports 10 - 20 as 802.1 authenticator ports.

```
HPswitch(config)# radius host 10.28.127.101 key rad4all
```

Configures the switch to look for a RADIUS server with an IP address of 10.28.127.101 and an encryption key of rad4all.

```
HPswitch(config)# aaa port-access authenticator e 10-20 unauth-vid 80
```

Configures ports 10 - 20 to use VLAN 80 as the Unauthorized-Client VLAN.

```
HPswitch(config)# aaa port-access authenticator e 10-20 auth-vid 81
```

Configures ports 10 - 20 to use VLAN 81 as the Authorized-Client VLAN.

```
HPswitch(config)# aaa port-access authenticator active
```

Activates 802.1x port-access on ports you have configured as authenticators.

Inspecting 802.1x Open VLAN Mode Operation. For information and an example on viewing current Open VLAN mode operation, refer to “Viewing 802.1x Open VLAN Mode Status” on page 47.

802.1x Open VLAN Operating Notes

- Although you can configure Open VLAN mode the same VLAN for both the Unauthorized-Client VLAN and the Authorized-Client VLAN, this is *not* recommended. Using the same VLAN for both purposes allows unauthenticated clients access to a VLAN intended only for authenticated clients, which poses a security breach.
- While an Unauthorized-Client VLAN is in use on a port, the switch temporarily removes the port from any other statically configured VLAN for which that port is configured as an untagged member. Note that the Menu interface will still display the port’s statically configured VLAN.
- An Unauthorized-Client VLAN should not be statically configured on any switch port that allows access to resources that must be protected from unauthenticated clients.
- If a port is configured as a tagged member of a VLAN that is not used as an Unauthorized-Client, Authorized-Client, or RADIUS-assigned VLAN, then the client can access such VLANs only if it is capable of operating in a tagged VLAN environment. Otherwise, the client can access only the Unauthorized-Client VLAN (before authentication) and either the Authorized-Client or RADIUS-assigned VLAN after authentication. (In all three cases, membership will be untagged, regardless of any static configuration specifying tagged membership.) If there is no Authorized-Client or RADIUS-assigned VLAN, then an authenticated client can access only a statically configured, untagged VLAN on that port.
- When a client’s authentication attempt on an Unauthorized-Client VLAN fails, the port remains a member of the Unauthorized-Client VLAN until the client disconnects from the port.
- During an authentication session on a port in 802.1x Open VLAN mode, if RADIUS specifies membership in an untagged VLAN, this assignment overrides port membership in the Authorized-Client VLAN. If there is no Authorized-Client VLAN configured, then the RADIUS assignment overrides any untagged VLAN for which the port is statically configured.
- If an authenticated client loses authentication during a session in 802.1x Open VLAN mode, the port VLAN membership reverts back to the Unauthorized-Client VLAN.

Option For Authenticator Ports: Configure Port-Security To Allow Only 802.1x Devices

If you are using port-security on authenticator ports, you can configure it to learn only the MAC address of the first 802.1x-aware device detected on the port. Then, only traffic from this specific device is allowed on the port. When this device logs off, another 802.1x-aware device can be authenticated on the port.

Syntax: port-security [ethernet] <port-list>

learn-mode port-access

Configures port-security on the specified port(s) to allow only the first 802.1x-aware device that the port detects.

action < none | send-alarm | send-disable >

Configures the port's response (in addition to blocking unauthorized traffic) to detecting an intruder.

Note

Port-Security operates with 802.1x authentication as described above only if the selected ports are configured as 802.1x; that is with the **control** mode in the port-access authenticator command set to **auto**. For example, to configure port 10 for 802.1x authenticator operation and display the result:

```
HPswitch(config)# aaa port-access authenticator e 10 control auto
HPswitch(config)# show port-access authenticator e 10 config
```

Note on Blocking a Non-802.1x Device

If the port's 802.1x authenticator **control** mode is configured to **authorized** (as shown below, instead of **auto**), then the first source MAC address from any device, whether 802.1x-aware or not, becomes the only authorized device on the port.

```
aaa port-access authenticator < port-list > control authorized
```

With 802.1x authentication disabled on a port or set to **authorized** (Force Authorize), the port may learn a MAC address that you don't want authorized. If this occurs, you can block access by the unauthorized, non-802.1x device by using one of the following options:

- If 802.1x authentication is disabled on the port, use these command syntaxes to enable it and allow only an 802.1x-aware device:

```
aaa port-access authenticator e < port-list >
```

Enables 802.1x authentication on the port.

```
aaa port-access authenticator e < port-list > control auto
```

Forces the port to accept only a device that supports 802.1x and supplies valid credentials.

If 802.1x authentication is enabled on the port, but set to **authorized** (Force Authorized), use this command syntax to allow only an 802.1x-aware device:

```
aaa port-access authenticator e < port-list > control auto
```

Forces the port to accept only a device that supports 802.1x and supplies valid credentials.

Configuring Switch Ports To Operate As Supplicants for 802.1x Connections to Other Switches

802.1x Authentication Commands	page 22
802.1x Supplicant Commands	
[no] aaa port-access < supplicant < [ethernet] < <i>port-list</i> >	page 42
[auth-timeout held-period start-period max-start initialize identity secret clear-statistics]	page 43
802.1x-Related Show Commands	page 45
RADIUS server configuration	pages 27

You can configure a switch port to operate as a supplicant in a connection to a port on another 802.1x-aware switch to provide security on links between 802.1x-aware switches. (Note that a port can operate as both an authenticator and a supplicant.)

For example, suppose that you want to connect two switches, where:

- Switch “A” has port 1 configured for 802.1x supplicant operation
- You want to connect port 1 on switch “A” to port 5 on switch “B”.

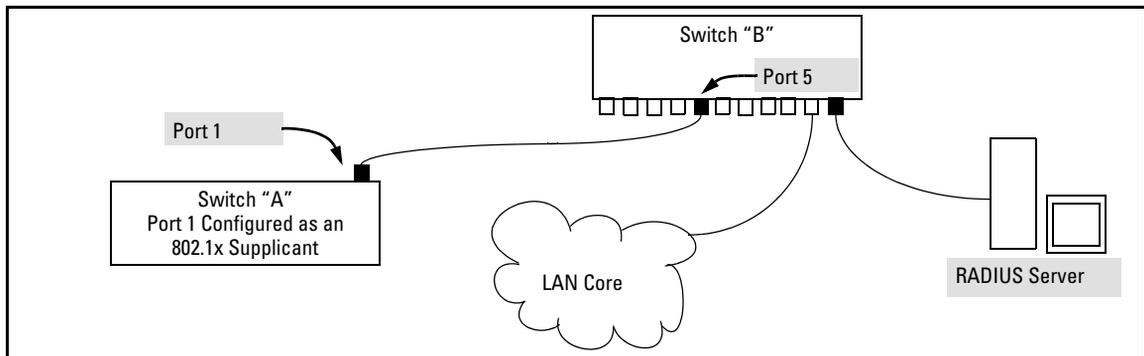


Figure 7. Example of Supplicant Operation

1. When port 1 on switch “A” is first connected to a port on switch “B”, or if the ports are already connected and either switch reboots, port 1 begins sending start packets to port 5 on switch “B”.

- If, after the supplicant port sends the configured number of start request packets, it does not receive a response, it assumes that switch “B” is not 802.1x-aware, and transitions to the authenticated state. If switch “B” is operating properly and is not 802.1x-aware, then the link should begin functioning normally, but without 802.1x security.
 - If, after sending one or more start request packets, port 1 receives a request packet from port 5, then switch “B” is operating as an 802.1x authenticator. The supplicant port then sends a response/ID packet. If switch “B” is configured for RADIUS authentication, it forwards this request to a RADIUS server. If switch “B” is configured for Local 802.1x authentication (page 26), the authenticator compares the switch “A” response to its local username and password.
2. The RADIUS server then responds with an access challenge that switch “B” forwards to port 1 on switch “A”.
 3. Port 1 replies with a hash response based on its unique credentials. Switch “B” forwards this response to the RADIUS server.
 4. The RADIUS server then analyzes the response and sends either a “success” or “failure” packet back through switch “B” to port 1.
 - A “success” response unblocks port 5 to normal traffic from port 1.
 - A “failure” response continues the block on port 5 and causes port 1 to wait for the “held-time” period before trying again to achieve authentication through port 5.

Note

You can configure a switch port to operate as both a supplicant and an authenticator at the same time.

Enabling a Switch Port To Operate as a Supplicant. You can configure one or more switch ports to operate as supplicants for point-to-point links to 802.1x-aware ports on other switches. *You must configure a port as a supplicant before you can configure any supplicant-related parameters.*

Syntax: [no] aaa port-access supplicant [ethernet] <port-list>

Configures a port to operate as a supplicant using either the default supplicant parameters or any previously configured supplicant parameters, whichever is the most recent. The “no” form of the command disables supplicant operation on the specified ports.

Configuring a Supplicant Switch Port. Note that you must enable supplicant operation on a port before you can change the supplicant configuration. This means you must execute the supplicant command once without any other parameters, then execute it again with a supplicant parameter you want to configure. If the intended authenticator port uses RADIUS authentication, then use the **identity** and **secret** options to configure the RADIUS-expected username and password on the supplicant port. If the intended authenticator port uses Local 802.1x authentication, then use the **identity** and **secret** options to configure the authenticator switch's local username and password on the supplicant port.

Syntax: aaa port-access supplicant [ethernet] < port-list >

*To enable supplicant operation on the designated ports, execute this command without any other parameters. After doing this, you can use the command again with the following parameters to configure supplicant operation. (Use one instance of the command for each parameter you want to configure. The **no** form disables supplicant operation on the designated port(s).*

[identity < username >]

Sets the username and password to pass to the authenticator port when a challenge-request packet is received from the authenticator port in response to an authentication request. If the intended authenticator port is configured for RADIUS authentication, then < username > and < password > must be the username and password expected by the RADIUS server. If the intended authenticator port is configured for Local authentication, then < username > and < password > must be the username and password configured on the Authenticator switch. (Defaults: Null)

[secret]

Enter secret: < password >

Repeat secret: < password >

Sets the secret password to be used by the port supplicant when an MD5 authentication request is received from an authenticator. The switch prompts you to enter the secret password after the command is invoked.

Syntax: aaa port-access supplicant [ethernet] < port-list > *(Syntax Continued)*

[auth-timeout < 1 - 300 >]

*Sets the period of time the port waits to receive a challenge from the authenticator. If the request times out, the port sends another authentication request, up to the number of attempts specified by the **max-start** parameter. (Default: 30 seconds).*

[max-start < 1 - 10 >]

Defines the maximum number of times the supplicant port requests authentication. See step 1 on page 41 for a description of how the port reacts to the authenticator response. (Default: 3).

[held-period < 0 - 65535 >]

Sets the time period the supplicant port waits after an active 802.1x session fails before trying to re-acquire the authenticator port. (Default: 60 seconds)

[start-period < 1 - 300 >]

*Sets the time period between Start packet retransmissions. That is, after a supplicant sends a start packet, it waits during the **start-period** for a response. If no response comes during the **start-period**, the supplicant sends a new start packet. The **max-start** setting (above) specifies how many start attempts are allowed in the session. (Default: 30 seconds)*

aaa port-access supplicant [ethernet] < port-list >

[initialize]

On the specified ports, blocks inbound and outbound traffic and restarts the 802.1x authentication process. Affects only ports configured as 802.1x supplicants.

[clear-statistics]

Clears and restarts the 802.1x supplicant statistics counters.

Displaying 802.1x Configuration, Statistics, and Counters

802.1x Authentication Commands	page 22
802.1x Supplicant Commands	page 41
802.1x Open VLAN Mode Commands	page 28
802.1x-Related Show Commands	
show port-access authenticator	below
show port-access supplicant	page 50
Details of 802.1x Mode Status Listings	page 47
RADIUS server configuration	pages 27

Show Commands for Port-Access Authenticator

Syntax: show port-access authenticator [[e] <port-list>]
 [config | statistics | session-counters]

- *Without* [<port-list> [config | statistics | session-counters]], *displays whether port-access authenticator is active (Yes or No) and the status of all ports configured for 802.1x authentication. The Authenticator Backend State in this data refers to the switch's interaction with the authentication server.*
- *With* <port-list> *only, same as above, but limits port status to only the specified port. Does not display data for a specified port that is not enabled as an authenticator.*
- *With* [<port-list> [config | statistics | session-counters]], *displays the [config | statistics | session-counters] data for the specified port(s). Does not display data for a specified port that is not enabled as an authenticator.*
- *With* [config | statistics | session-counters] *only, displays the [config | statistics | session-counters] data for all ports enabled as authenticators.*

For descriptions of [config | statistics | session-counters] refer to the next section of this table.

Syntax: show port-access authenticator (*Syntax Continued*)

config [[e] < port-list >]

Shows:

- *Whether port-access authenticator is active*
- *The 802.1x configuration of the ports configured as 802.1x authenticators*

If you do not specify < port-list >, the command lists all ports configured as 802.1x port-access authenticators. Does not display data for a specified port that is not enabled as an authenticator.

statistics [[e] < port-list >]

Shows:

- *Whether port-access authenticator is active*
- *The statistics of the ports configured as 802.1x authenticators, including the supplicant's MAC address, as determined by the content of the last EAPOL frame received on the port.*

Does not display data for a specified port that is not enabled as an authenticator.

session-counters [[e] < port-list >]

Shows:

- *Whether port-access authenticator is active*
- *The session status on the specified ports configured as 802.1x authenticators*

*Also, for each port, the “User” column lists the user name the supplicant included in its response packet. (For the switch, this is the **identity** setting included in the **supplicant** command—page 43.) Does not display data for a specified port that is not enabled as an authenticator.*

Viewing 802.1x Open VLAN Mode Status

You can examine the switch's current VLAN status by using the **show port-access authenticator** and **show vlan <vlan-id>** commands as illustrated in this section. Figure 8 shows an example of **show port-access authenticator** output, and table 3 describes the data that this command displays. Figure 9 shows related VLAN data that can help you to see how the switch is using statically configured VLANs to support 802.1x operation.

```

HPswitch(config)# show port-access authenticator 1-4

```

Port Access Authenticator Status

Port-access authenticator activated [No] : Yes

Port	Status	Access Control	Authenticator State	Authenticator Backend State	Unauth VLAN ID	Auth VLAN ID	Current VLAN ID
1	Closed	Auto	Connecting	Idle	(100)	101	(100)
2	(Open)	Auto	(Authorized)	Idle	100	(101)	(101)
3	Closed	Auto	Connecting	Idle	100	(0)	100
4	Closed	Auto	Disconnected	Idle	100	101	(No PVID)

An Unauth VLAN ID appearing in the Current VLAN ID column for the same port indicates an unauthenticated client is connected to this port.
(Assumes that the port is not a statically configured member of VLAN 100.)

Items 1 through 3 indicate that an authenticated client is connected to port 2:

1. **Open** in the Status column
2. **Authorized** in the Authenticator State column
3. The Auth VLAN ID (**101**) is also in the Current VLAN ID column. (This assumes that the port is not a statically configured member of VLAN 101.)

4. A "0" in the row for port 3 indicates there is no Authorized VLAN configured for port 3.

5. "No PVID" means there is currently no untagged VLAN membership on port 4.

Figure 8. Example Showing Ports Configured for Open VLAN Mode

Thus, in the **show port-access authenticator** output:

- When the **Auth VLAN ID** is configured and matches the **Current VLAN ID** in the above command output, an authenticated client is connected to the port. (This assumes the port is not a statically configured member of the VLAN you are using for Auth VLAN.)
- When the **Unauth VLAN ID** is configured and matches the **Current VLAN ID** in the above command output, an unauthenticated client is connected to the port. (This assumes the port is not a statically configured member of the VLAN you are using for Unauth VLAN.)

Note that because a temporary Open VLAN port assignment to either an authorized or unauthorized VLAN is an untagged VLAN membership, these assignments temporarily replace any other untagged VLAN membership that is statically configured on the port. For example, if port 12 is statically configured as an untagged member of VLAN 1, but is configured to use VLAN 25 as an authorized VLAN, then the port's membership in VLAN 1 will be temporarily suspended whenever an authenticated 802.1x client is attached to the port.

Table 4. Open VLAN Mode Status

Status Indicator	Meaning
Port	Lists the ports configured as 802.1x port-access authenticators.
Status	<p>Closed: Either no client is connected or the connected client has not received authorization through 802.1x authentication.</p> <p>Open: An authorized 802.1x supplicant is connected to the port.</p>
Access Control	
This state is controlled by the following port-access command syntax:	
<p>HPswitch(config)# aaa port-access authenticator < port-list > control < authorized auto unauthorized ></p>	
<p>Auto: Configures the port to allow network access to any connected device that supports 802.1x authentication and provides valid 802.1x credentials. (This is the default authenticator setting.)</p> <p>FA: Configures the port for "Force Authorized", which allows access to any device connected to the port, regardless of whether it meets 802.1x criteria. (You can still configure console, Telnet, or SSH security on the port.)</p> <p>FU: Configures the port for "Force Unauthorized", which blocks access to any device connected to the port, regardless of whether the device meets 802.1x criteria.</p>	
Authenticator State	<p>Connecting: A client is connected to the port, but has not received 802.1x authentication.</p> <p>Force Unauth: Indicates the "Force Unauthorized" state. Blocks access to the network, regardless of whether the client supports 802.1x authentication or provides 802.1x credentials.</p> <p>Force Auth: Indicates the "Force Authorized" state. Grants access to any device connected to the port. The device does not have to support 802.1x authentication or provide 802.1x credentials.</p> <p>Authorized: The device connected to the port supports 802.1x authentication, has provided 802.1x credentials, and has received access to the network. This is the default state for access control.</p> <p>Disconnected: No client is connected to the port.</p>
Authenticator Backend State	<p>Idle: The switch is not currently interacting with the RADIUS authentication server. Other states (Request, Response, Success, Fail, Timeout, and Initialize) may appear temporarily to indicate interaction with a RADIUS server. However, these interactions occur quickly and are replaced by Idle when completed.</p>
Unauthorized VLAN ID	<p>< vlan-id >: Lists the VID of the static VLAN configured as the unauthorized VLAN for the indicated port.</p> <p>0: No unauthorized VLAN has been configured for the indicated port.</p>
Authorized VLAN ID	<p>< vlan-id >: Lists the VID of the static VLAN configured as the authorized VLAN for the indicated port.</p> <p>0: No authorized VLAN has been configured for the indicated port.</p>

Status Indicator	Meaning
Current VLAN ID	< vlan-id >: Lists the VID of the static, untagged VLAN to which the port currently belongs.
No PVID:	The port is not an untagged member of any VLAN.

Syntax: show vlan < vlan-id >

Displays the port status for the selected VLAN, including an indication of which port memberships have been temporarily overridden by Open VLAN mode.

```

HPswitch(config)# show vlan 1
Status and Counters - VLAN Information - Ports - VLAN 1
802.1Q VLAN ID : 1
Name           : DEFAULT_VLAN
Status         : Static

Port Information Mode      Unknown VLAN Status
-----
1              Untagged Learn      Up
2              Untagged Learn      Up
3              Untagged Learn      Up
4              Untagged Learn      Up
12             Untagged Learn      Up
14             Tagged Learn       Up
5              Untagged Learn      Down
:              :                :
:              :                :
23             B              Untagged Learn      Up
24             B              Untagged Learn      Up

Overridden Port VLAN configuration

Port Mode
-----
1      Untagged
3      Untagged

```

Note that ports 1 and 3 are not in the upper listing, but are included under "Overridden Port VLAN configuration". This shows that static, untagged VLAN memberships on ports 1 and 3 have been overridden by temporary assignment to the authorized or unauthorized VLAN. Using the **show port-access authenticator < port-list >** command shown in figure 8 provides details.

Figure 9. Example of Showing a VLAN with Ports Configured for Open VLAN Mode

Show Commands for Port-Access Supplicant

Syntax: show port-access supplicant [[e] <port-list>] [statistics]

show port-access supplicant [[e] <port-list>]

*Shows the port-access supplicant configuration (excluding the **secret** parameter) for all ports or <port-list> ports configured on the switch as supplicants. The Supplicant State can include the following:*

Connecting - Starting authentication.

Authenticated - Authentication completed (regardless of whether the attempt was successful).

Acquired - The port received a request for identification from an authenticator.

Authenticating - Authentication is in progress.

Held - Authenticator sent notice of failure. The supplicant port is waiting for the authenticator's held-period (page 43).

For descriptions of the supplicant parameters, refer to “Configuring a Supplicant Switch Port” on page 43.

show port-access supplicant [[e] <port-list>] statistics

Shows the port-access statistics and source MAC address(es) for all ports or <port-list> ports configured on the switch as supplicants. See the “Note on Supplicant Statistics”, below.

Note on Supplicant Statistics. For each port configured as a supplicant, **show port-access supplicant statistics [e] <port-list>** displays the source MAC address and statistics for transactions with the authenticator device most recently detected on the port. If the link between the supplicant port and the authenticator device fails, the supplicant port continues to show data received from the connection to the most recent authenticator device until one of the following occurs:

- The supplicant port detects a different authenticator device.
- You use the **aaa port-access supplicant [e] <port-list> clear-statistics** command to clear the statistics for the supplicant port.
- The switch reboots.

Thus, if the supplicant's link to the authenticator fails, the supplicant retains the transaction statistics it most recently received until one of the above events occurs. Also, if you move a link with an authenticator from one supplicant port to another without clearing the statistics data from the first port, the authenticator's MAC address will appear in the supplicant statistics for both ports.

How RADIUS/802.1x Authentication Affects VLAN Operation

Static VLAN Requirement. RADIUS authentication for an 802.1x client on a given port can include a (static) VLAN requirement. (Refer to the documentation provided with your RADIUS application.) The static VLAN to which a RADIUS server assigns a client must already exist on the switch. If it does not exist or is a dynamic VLAN (created by GVRP), authentication fails. Also, for the session to proceed, the port must be an untagged member of the required VLAN. If it is not, the switch temporarily reassigns the port as described below.

If the Port Used by the Client Is Not Configured as an Untagged Member of the Required Static VLAN: When a client is authenticated on port “N”, if port “N” is not already configured as an untagged member of the static VLAN specified by the RADIUS server, then the switch temporarily assigns port “N” as an untagged member of the required VLAN (for the duration of the 802.1x session). *At the same time, if port “N” is already configured as an untagged member of another VLAN, port “N” loses access to that other VLAN for the duration of the session.* (This is because a port can be an untagged member of only one VLAN at a time.)

For example, suppose that a RADIUS-authenticated, 802.1x-aware client on port 2 requires access to VLAN 22, but VLAN 22 is configured for no access on port 2, and VLAN 33 is configured as untagged on port 2:

```

===== CONSOLE - MANAGER MODE =====
Switch Configuration - VLAN - VLAN Port Assignment

Port  default_vlan  vlan_22  vlan_33  vlan_44
----+-----
1  | Untagged  Tagged   No       No
2  | No        No       Untagged No
3  | Untagged  Forbid  Forbid   Forbid
4  | Untagged  Tagged  Tagged   Tagged
:   |         :       :       :
:   |         :       :       :
Actions->  Cancel  Edit    Save    Help
Cancel changes and return to previous screen.
Use arrow keys to change action selection and <Enter> to execute.
    
```

Scenario: An authorized 802.1x client requires access to VLAN 22 from port 2. However, access to VLAN 22 is blocked (not untagged or tagged) on port 2 and VLAN 33 is untagged on port 2.

Figure 10. Example of an Active VLAN Configuration

In figure 10, if RADIUS authorizes an 802.1x client on port 2 with the requirement that the client use VLAN 22, then:

- VLAN 22 becomes available as Untagged on port 2 for the duration of the session.

Enhancements in Release F.05.xx

- VLAN 33 becomes unavailable to port 2 for the duration of the session (because there can be only one untagged VLAN on any port).

You can use the **show vlan** <vlan-id> command to view this temporary change to the active configuration, as shown below:

- You can see the temporary VLAN assignment by using the **show vlan** <vlan-id> command with the <vlan-id> of the static VLAN that the authenticated client is using.

```
HPswitch(config)# show vlan 22
Status and Counters - VLAN Information - Ports - VLAN 22
802.1Q VLAN ID : 22
Name           : vlan_22
Status         : Static

Port Information Mode      Unknown VLAN Status
-----
1                Tagged    Learn          Up
2                (802.1X) Learn          Up
4                Tagged    Learn          Up
.                .          .              .
.                .          .              .
.                .          .              .

Overridden Port VLAN configuration
Port  Mode
----  ---
2     (No)

This entry shows that port 2 is temporarily untagged on
VLAN 22 for an 802.1x session. This is to accommodate an
802.1x client's access, authenticated by a RADIUS server,
where the server included an instruction to put the client's
access on VLAN 22.

Note: With the current VLAN configuration (figure 10), the
only time port 2 appears in this show vlan 22 listing is during
an 802.1x session with an attached client. Otherwise, port
2 is not listed.
```

Figure 11. The Active Configuration for VLAN 22 Temporarily Changes for the 802.1x Session

- With the preceding in mind, since (static) VLAN 33 is configured as untagged on port 2 (see figure 10), and since a port can be untagged on only one VLAN, port 2 loses access to VLAN 33 for the duration of the 802.1x session involving VLAN 22. You can verify the temporary loss of access to VLAN 33 with the **show vlan 33** command.

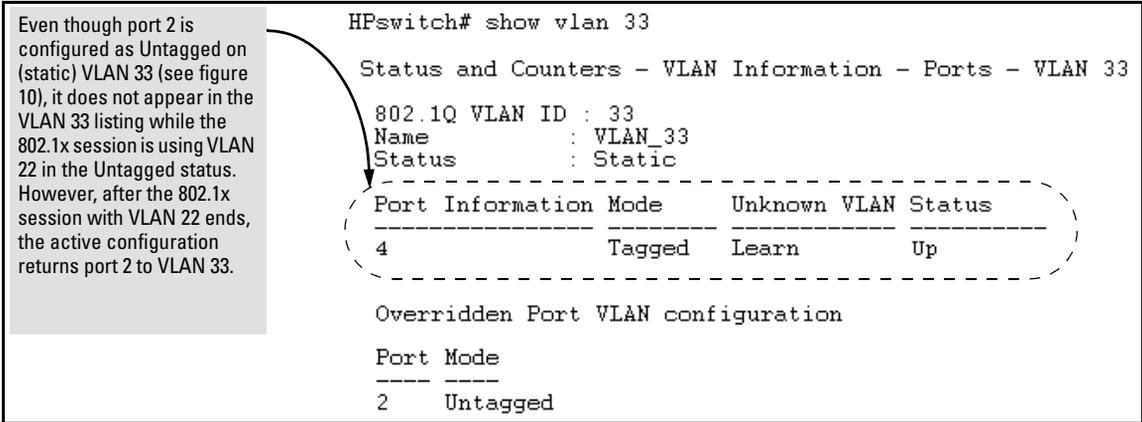


Figure 12. The Active Configuration for VLAN 33 Temporarily Drops Port 22 for the 802.1x Session

When the 802.1x client’s session on port 2 ends, the port discards the temporary untagged VLAN membership. At this time the static VLAN actually configured as untagged on the port again becomes available. Thus, when the RADIUS-authenticated 802.1x session on port 2 ends, VLAN 22 access on port 2 also ends, and the untagged VLAN 33 access on port 2 is restored.

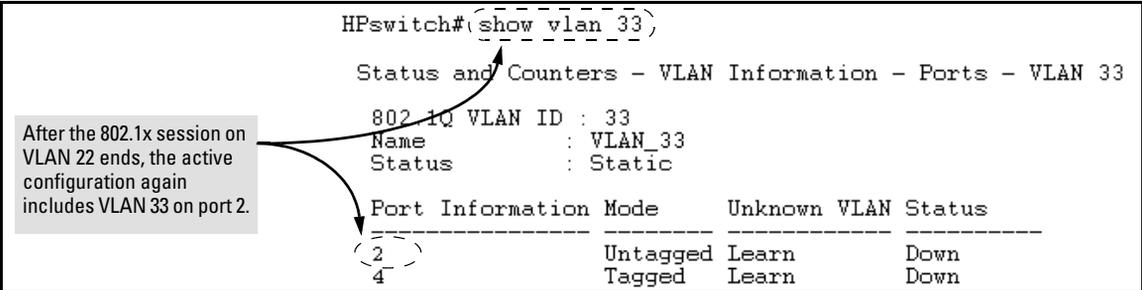


Figure 13. The Active Configuration for VLAN 33 Restores Port 2 After the 802.1x Session Ends

Notes

Any port VLAN-ID changes you make on 802.1x-aware ports during an 802.1x-authenticated session do not take effect until the session ends.

With GVRP enabled, a temporary, untagged static VLAN assignment created on a port by 802.1x authentication is advertised as an existing VLAN. If this temporary VLAN assignment causes the switch to disable a configured (untagged) static VLAN assignment on the port, then the disabled VLAN assignment is not advertised. When the 802.1x session ends, the switch:

- Eliminates and ceases to advertise the temporary VLAN assignment.
 - Re-activates and resumes advertising the temporarily disabled VLAN assignment.
-

Messages Related to 802.1x Operation

Table 5. 802.1x Operating Messages

Message	Meaning
Port < port-list > is not an authenticator.	The ports in the port list have not been enabled as 802.1x authenticators. Use this command to enable the ports as authenticators: HPswitch(config)# aaa port-access authenticator e 10
Port < port-list > is not a supplicant.	Occurs when there is an attempt to change the supplicant configuration on a port that is not currently enabled as a supplicant. Enable the port as a supplicant and then make the desired supplicant configuration changes. Refer to "Enabling a Switch Port To Operate as a Supplicant" on page 42.
No server(s) responding.	This message can appear if you configured the switch for EAP-RADIUS or CHAP-RADIUS authentication, but the switch does not receive a response from a RADIUS server. Ensure that the switch is configured to access at least one RADIUS server. (Use show radius .) If you also see the message Can't reach RADIUS server < x.x.x.x >, try the suggestions listed for that message (page 114).

Message	Meaning
LACP has been disabled on 802.1x port (s) .	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(s), and enables 802.1x on that port.
Error configuring port < port-number >: LACP and 802.1x cannot be run together.	Also, the switch will not allow you to configure LACP on a port on which port access (802.1x) is enabled.

IGMP Version 3 Support

When the switch receives an IGMPv3 Join, it accepts the host request and begins forwarding the IGMP traffic. This means that ports that have not joined the group and are not connected to routers or the IGMP Querier will not receive the group's multicast traffic.

The switch does not support the IGMPv3 "Exclude Source" or "Include Source" options in the Join Reports. Instead, the group is simply joined from all sources.

The switch does not support becoming a version 3 Querier. It will become a version 2 Querier in the absence of any other Querier on the network.

— *This page is intentionally unused.* —

Enhancements in Release F.04.08

Enhancement	Summary	Page
Friendly Port Names	Enables you to assign optional, meaningful names to physical ports on the switch.	58
Security Enhancements		
SSH Security	Provide remote access to management functions on the switches via encrypted paths between the switch and management station clients capable of SSHv1 operation.	63
RADIUS	Protect access to the switch and monitor use of network resources through a centralized client authentication and accounting service.	89
Port-Based Access Control (802.1x)	Release F.04.08 provides 802.1x port-access control for users requesting access from multiple points within the network, including application of user profiles configured on a central RADIUS server. Release F.05.17 updates this feature to include Open VLAN mode, which changes how the port responds when it detects a new client. For this reason you will find the documentation for the Port-Access (802.1x) with Open VLAN mode under "Enhancements in Release F.05.17" instead of in this section.	14
IP Preserve	Enable retention of the current IP address and subnet mask (for the switch's default VLAN), and the default gateway address when downloading a configuration file and rebooting the switch. (Operates on switches that use the Manual IP addressing instead of the default DHCP method.)	117
QoS Priority	Enable assignment of non-default priority settings to inbound, untagged packets received on the switch.	120
Isolated Port Groups	Provides an alternative to VLANs in situations where VLANs cannot be used. Release F.05.17 updates this feature to include two new groups. For this reason you will find the documentation for Isolated Port Groups under "Enhancements in Release F.05.17" instead of in this section.	6
Terminating Remote Sessions	Provides a "kill" command to terminate remote Telnet and SSH sessions.	124
Rapid Spanning-Tree (802.1W) (RSTP)	Provides the functionality for the new Spanning Tree standard, IEEE 802.1w (RSTP), which is supported by the G.04.04 (or greater) release of your switch software	125

Using Friendly (Optional) Port Names

Feature	Default	Menu	CLI	Web
Configure Friendly Port Names	Standard Port Numbering	n/a	page 59	n/a
Display Friendly Port Names	n/a	n/a	page 60	n/a

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 any port on the switch. 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” on page 60.)
- 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

Syntax: interface [e] <port-list> name <port-name-string> *Assigns a port name to port-list.*
 no interface [e] <port-list> name *Deletes the port name from port-list.*

Configuring a Single Port Name. Suppose that you have connected port 3 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 3:

```
HPswitch(config)# int e 3 name Bill_Smith@10.25.101.73
HPswitch(config)# write mem
HPswitch(config)# show name 3
  Port Names
  Port : 3
    Type : 10/100TX
    Name : Bill_Smith@10.25.101.73
```

Figure 14. Example of Configuring a Friendly Port Name

Configuring the Same Name for Multiple Ports. Suppose that you want to use ports 5 through 8 as a trunked link to a server used by a drafting group. In this case you might configure ports 5 through 8 with the name "Draft-Server:Trunk".

```
HPswitch(config)# int e 5-8 name Draft-Server:Trunk
HPswitch(config)# write mem
HPswitch(config)# show name 5-8
  Port Names

  Port : 5
    Type : 10/100TX
    Name : Draft-Server:Trunk

  Port : 6
    Type : 10/100TX
    Name : Draft-Server:Trunk

  Port : 7
    Type : 10/100TX
    Name : Draft-Server:Trunk

  Port : 8
    Type : 10/100TX
    Name : Draft-Server:Trunk
```

Figure 15. Example of Configuring One Friendly Port Name on Multiple Ports

Displaying Friendly Port Names with Other Port Data

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

- **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.)
- **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.)
- **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.)

To List All Ports or Selected Ports with Their Friendly Port Names. This command lists names assigned to a specific port.

Syntax: show name [port-list] *Lists the friendly port name with its corresponding port number and port type. show name alone lists this data for all ports on the switch.*

For example:

```

HPswitch(config)# show name
Port Names
Port  Type      Name
-----
 1   10/100TX  not assigned
 2   10/100TX  not assigned
 3   10/100TX  [ Bill_Smith@10.25.101.73 ]
 4   10/100TX  not assigned
 5   10/100TX  [ Draft-Server:Trunk ]
 6   10/100TX  | Draft-Server:Trunk |
 7   10/100TX  | Draft-Server:Trunk |
 8   10/100TX  | Draft-Server:Trunk |
 9   10/100TX  not assigned
10   10/100TX  not assigned
11   10/100TX  not assigned
12   10/100TX  not assigned
:   :         :
:   :         :

```

Figure 16. Example of Friendly Port Name Data for All Ports on the Switch

```

HPswitch(config)# show name 2,3,5
Port Names
┌───┬───┬───┬───┬───┬───┐
│ Port : 2 │   │   │   │   │   │
│ Type : 10/100TX │   │   │   │   │   │
│ Name : not assigned │   │   │   │   │   │
└───┴───┴───┴───┴───┴───┘
Port : 3
Type : 10/100TX
Name : Bill_Smith@10.25.101.73
Port : 5
Type : 10/100TX
Name : Draft-Server:Trunk

```

Port Without a "Friendly" Name

Friendly port names assigned in previous examples.

Figure 17. Example of Friendly Port Name Data for Specific Ports on the Switch

Including Friendly Port Names in Per-Port Statistics Listings. A friendly port name configured to a port is automatically included when you display the port's statistics output.

Syntax: show interface <port-number> *Includes the friendly port name with the port's traffic statistics listing.*

For example, if you configure port 1 with the name "O'Connor_10.25.101.43", the show interface output for this port appears similar to the following:

```

HPswitch(config)# show interface 1
Status and Counters - Port Counters for port 1

Name   : O'Connor@10.25.101.43
Link Status      : Up

Bytes Rx       : 894,568           Bytes Tx       : 2470
Unicast Rx     : 1179             Unicast Tx     : 13
Bcast/Mcast Rx : 5280             Bcast/Mcast Tx : 13

FCS Rx        : 36                Drops Tx       : 0
Alignment Rx  : 2                 Collisions Tx  : 0
Runts Rx      : 0                 Late Colln Tx : 0
Giants Rx     : 0                 Excessive Colln : 0
Total Rx Errors : 38              Deferred Tx    : 0

```

Friendly Port Name

Figure 18. Example of a Friendly Port Name in a Per-Port Statistics Listing

Enhancements in Release F.04.08

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
```

To Search the Configuration for Ports with Friendly Port Names. This option tells you which friendly port names have been saved to the startup-config file. (The **show config** command 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.*

For example, if you configure port 1 with a friendly port name:

```
HPswitch(config)# int e 1 name Print_Server@10.25.101.43
HPswitch(config)# write mem
HPswitch(config)# int e 2 name Herbert's_PC

HPswitch(config)# show config

Startup configuration:
; J4865A Configuration Editor; Created on release #F.05.17
hostname "HPswitch"
time daylight-time-rule None
no cdp run
interface 1
  name "Print Server@10.25.101.43"
exit
snmp-server community "public" Unrestricted
vlan 1
  name "DEFAULT_VLAN"
  untagged 1-14
  ip address dhcp-bootp
  exit
no aaa port-access authenticator active
```

This command sequence saves the friendly port name for port 1 in the startup-config file, but does not do so for the name entered for port 2.

Listing includes friendly port name for port 1 only.

In this case, **show config** lists only port 1. Executing **write mem** after entering the name for port 2, and then executing **show config** again would result in a listing that includes both ports.

Figure 19. Example Listing of the Startup-Config File with a Friendly Port Name Configured (and Saved)

Configuring Secure Shell (SSH)

Feature	Default	Menu	CLI	Web
Generating a public/private key pair on the switch	No	n/a	page 70	n/a
Using the switch's public key	n/a	n/a	page 72	n/a
Enabling SSH	Disabled	n/a	page 74	n/a
Enabling client public-key authentication	Disabled	n/a	pages 77, 80	n/a
Enabling user authentication	Disabled	n/a	page 77	n/a

The Series 2500 switches use Secure Shell version 1 (SSHv1) to provide remote access to management functions on the switches via encrypted paths between the switch and management station clients capable of SSHv1 operation. (The switches can be authenticated by SSHv2 clients that support SSHv1.) However, to use the reverse option—authenticating an SSHv2 user to the switch—you must have a method for converting the SSHv2 PEM public-key format to non-encoded ASCII. Refer to "PEM (Privacy Enhanced Mode)" on page 65.

SSH provides Telnet-like functions but, unlike Telnet, SSH provides encrypted, authenticated transactions. The authentication types include:

- Client public-key authentication
- Switch SSH and user password authentication

Client Public Key Authentication (Login/Operator Level) with User Password Authentication (Enable/Manager Level). This option uses one or more public keys (from clients) that must be stored on the switch. Only a client with a private key that matches a stored public key can gain access to the switch. (The same private key can be stored on one or more clients.)

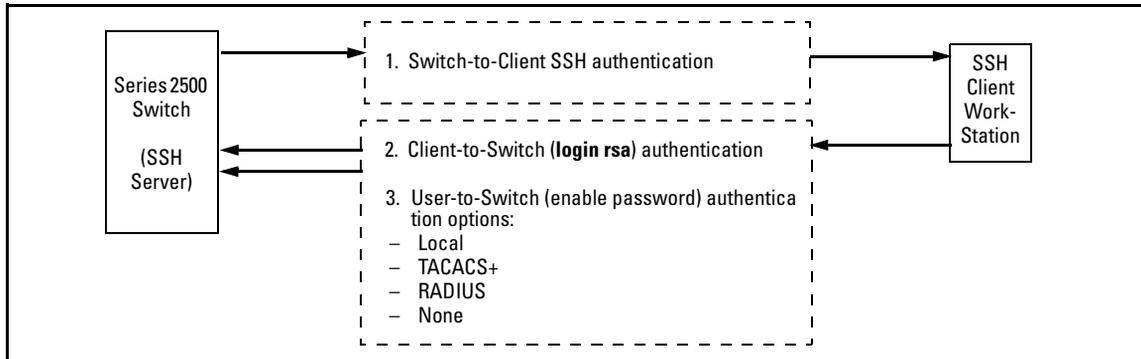


Figure 20. Client Public Key Authentication Model

Note

SSH in the HP ProCurve Series 2500 switches is based on the OpenSSH software toolkit. For more information on OpenSSH, visit <http://www.openssh.com>.

Switch SSH and User Password Authentication . This option is a subset of the client public-key authentication show in figure 20. It occurs if the switch has SSH enabled but does not have login access (**login rsa**) configured to authenticate the client's key. As in figure 20, the switch authenticates itself to SSH clients. Users on SSH clients then authenticate themselves to the switch (login and/or enable levels) by providing passwords stored locally on the switch or on a TACACS+ or RADIUS server. However, the client does not use a key to authenticate itself to the switch.

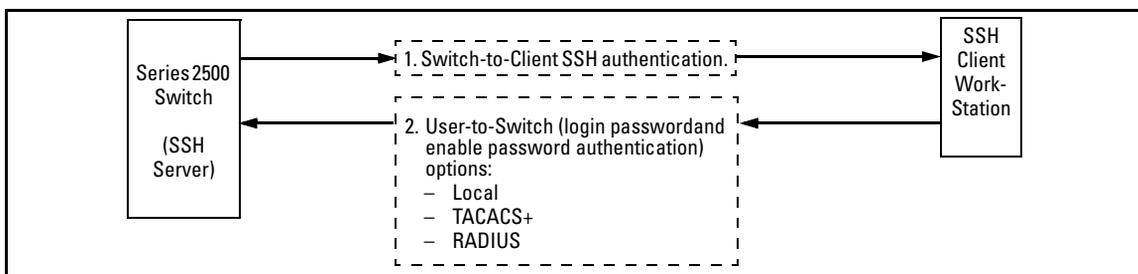


Figure 21. Switch/User Authentication

SSH on the Series 2500 switches supports these data encryption methods:

- 3DES (168-bit)
- DES (56-bit)

Note

This release supports SSH version 1 only, and all references to SSH in this document are to SSHv1 unless otherwise stated. SSH version 1 uses RSA public key algorithms exclusively, and all references to either a public or private key mean keys generated using these algorithms unless otherwise noted.

Terminology

- **SSH Server:** An HP Series 2500 switch with SSH enabled.
- **Key Pair:** A pair of keys generated by the switch or an SSH client application. Each pair includes a public key (that can be read by anyone) and a private key that is held internally in the switch or by a client.
- **PEM (Privacy Enhanced Mode):** Refers to an ASCII-formatted client public-key that has been encoded for greater security. SSHv2 client public-keys are typically stored in the PEM format. See figures 22 and 23 for examples of PEM-encoded ASCII and non-encoded ASCII keys.
- **Private Key:** An internally generated key used in the authentication process. A private key generated by the switch is not accessible for viewing or copying. A private key generated by an SSH client application is typically stored in a file on the client device and, together with its public key counterpart, can be copied and stored on multiple devices.
- **Public Key:** An internally generated counterpart to a private key. Public keys are used for authenticating a
- **Enable Level:** Manager privileges on the switch.
- **Login Level:** Operator privileges on the switch.
- **Local password or username:** A Manager-level or Operator-level password configured in the switch.
- **SSH Enabled:** (1) A public/private key pair has been generated on the switch (**crypto key generate [rsa]**) and (2) SSH is enabled (**ip ssh**). (You can generate a key pair without enabling SSH, but you cannot enable SSH without first generating a key pair. See “2. Generating the Switch’s Public and Private Key Pair” on page 70 and “4. Enabling SSH on the Switch and Anticipating SSH Client Contact Behavior” on page 74.)

Prerequisite for Using SSH

Before using a Series 2500 switch as an SSH server, you must install a publicly or commercially available SSH client application on the computer(s) you use for management access to the switch. If you want client public-key authentication (page 63), then the client program must have the capability to generate public and private key pairs.

Public Key Format Requirement

Any client application you use for client public-key authentication with the switch must have the capability to store a public key in non-encoded ASCII format. The switch does not interpret keys generated using the PEM (Privacy Enhanced Mode) format (also in ASCII characters) that some SSHv2 client applications use for storing public keys. If your client application stores PEM-encoded

keys by default, check the application software for a key conversion utility or use a third-party key conversion utility.

```

"Pub Key Gen 21 Dec 2001 12:01"A1B3Nz1y2+orEML . . . Q8D8qDM1ozu1c="*** End of Pub Key ***"
    
```

Figure 22. Example of Public Key in PEM-Encoded ASCII Format Common for SSHv2 Clients

```

512 37 78193303392019545793321845914508115859448079486918367079008218589443776362026267. . .
    
```

Figure 23. Example of Public Key in Non-Encoded ASCII Format (Common for SSHv1 Client Applications)

Steps for Configuring and Using SSH for Switch and Client Authentication

For two-way authentication between the switch and an SSH client, you must use the login (Operator) level.

Table 6. SSH Options

Switch Access Level	Primary SSH Authentication	Authenticate Switch Public Key to SSH Clients?	Authenticate Client Public Key to the Switch?	Primary Switch Password Authentication	Secondary Switch Password Authentication
Operator (Login) Level	ssh login rsa	Yes	Yes ¹	No ¹	local or none
	ssh login Local	Yes	No	Yes	local or none
	ssh login TACACS	Yes	No	Yes	local or none
	ssh login RADIUS	Yes	No	Yes	local or none
Manager (Enable) Level	ssh enable local	Yes	No	Yes	local or none
	ssh enable tacacs	Yes	No	Yes	local or none
	ssh enable radius	Yes	No	Yes	local or none

¹ For **ssh login rsa**, the switch uses client public-key authentication instead of the switch password options for primary authentication.

The general steps for configuring SSH include:

A. Client Preparation

1. Install an SSH client application on a management station you want to use for access to the switch. (Refer to the documentation provided with your SSH client application.)
2. Optional—If you want the switch to authenticate a client public-key on the client:
 - a. Either generate a public/private key pair on the client computer or (if your client application allows) or import a client key pair that you have generated using another SSH application.
 - b. Copy the client public key into an ASCII file on a TFTP server accessible to the switch and download the client public key file to the switch. (The client public key file can hold up to 10 client keys.) This topic is covered under “To Create a Client-Public-Key Text File” on page 81.

B. Switch Preparation

1. Assign a login (Operator) and enable (Manager) password on the switch (page 70).
2. Generate a public/private key pair on the switch (page 70).
You need to do this only once. The key remains in the switch even if you reset the switch to its factory-default configuration. (You can remove or replace this key pair, if necessary.)
3. Copy the switch’s public key to the SSH clients you want to access the switch (page 72).
4. Enable SSH on the switch (page 74).
5. Configure the primary and secondary authentication methods you want the switch to use. In all cases, the switch will use its host-public-key to authenticate itself when initiating an SSH session with a client.

•SSH Login (Operator) options:

–Option A:

Primary: Local, TACACS+, or RADIUS password
Secondary: Local password or none

–Option B:

Primary: Client public-key authentication (**login rsa** — page 80)
Secondary: Local password or none

Note that if you want the switch to perform client public-key authentication, you must configure the switch with Option B.

•SSH Enable (Manager) options:

Primary: Local, TACACS+, or RADIUS
Secondary: Local password or none

6. Use your SSH client to access the switch using the switch's IP address or DNS name (if allowed by your SSH client application). Refer to the documentation provided with the client application.

General Operating Rules and Notes

- Any SSH client application you use must offer backwards-compatibility to SSHv1 keys and operation.
- Public keys generated on an SSH client computer must be in ASCII format (used in SSHv1) if you want to be able to authenticate a client to the switch. The switch does not support keys generated in the PEM (base-64 Privacy Enhanced Mode) format. See the Note under "Prerequisite for Using SSH" on page 65.
- The switch's own public/private key pair and the (optional) client public key file are stored in the switch's flash memory and are not affected by reboots or the **erase startup-config** command.
- Once you generate a key pair on the switch you should avoid re-generating the key pair without a compelling reason. Otherwise, you will have to re-introduce the switch's public key on all management stations (clients) you previously set up for SSH access to the switch. In some situations this can temporarily allow security breaches.
- When stacking is enabled, SSH provides security only between an SSH client and the stack manager. Communications between the stack commander and stack members is not secure.
- The switch does not support outbound SSH sessions. Thus, if you Telnet from an SSH-secure switch to another SSH-secure switch, *the session is not secure*.

Configuring the Switch for SSH Operation

SSH-Related Commands in This Section

show ip ssh	page 76
show ip client-public-key [< babble fingerprint >]	page 83
show ip host-public-key [< babble fingerprint >]	page 73
show authentication	page 79
crypto key < generate zeroize > [rsa]	page 71
ip ssh	page 75
key-size < 512 768 1024 >	page 75
port < 1 - 65535 >	page 75
timeout < 5 .. 120 >	page 75
aaa authentication ssh	
login < local tacacs radius rsa >	page 77, 78
< local none >	page 77
enable < tacacs radius local >	page 77
< local none >	page 77
copy tftp pub-key-file <tftp server IP> <public key file>	page 83
clear public key	page 83

1. Assigning a Local Login (Operator) and Enable (Manager) Password

At a minimum, HP recommends that you always assign at least a Manager password to the switch. Otherwise, under some circumstances, anyone with Telnet, web, or serial port access could modify the switch's configuration.

To Configure Local Passwords. You can configure both the Operator and Manager password with one command.

Syntax: password < manager | operator | all >

```
HPswitch(config)# password all
New password for Operator: *****
Please retype new password for Operator: *****
New password for Manager: *****
Please retype new password for Manager: *****
HPswitch(config)#
```

Figure 24. Example of Configuring Local Passwords

2. Generating the Switch's Public and Private Key Pair

You must generate a public and private host key pair on the switch. The switch uses this key pair, along with a dynamically generated session key pair to negotiate an encryption method and session with an SSH client trying to connect to the switch.

The host key pair is stored in the switch's flash memory, and only the public key in this pair is readable. The public key should be added to a "known hosts" file (for example, \$HOME/.ssh/known_hosts on UNIX systems) on the SSH clients who you want to have access to the switch. Some SSH client applications automatically add the the switch's public key to a "known hosts" file. Other SSH applications require you to manually create a known hosts file and place the switch's public key in the file. (Refer to the documentation for your SSH client application.)

(The session key pair mentioned above is not visible on the switch. It is a temporary, internally generated pair used for a particular switch/client session, and then discarded.)

Notes

When you generate a host key pair on the switch, the switch places the key pair in flash memory (and not in the running-config file). Also, the switch maintains the key pair across reboots, including power cycles. You should consider this key pair to be "permanent"; that is, avoid re-generating the key pair without a compelling reason. Otherwise, you will have to re-introduce the switch's public key on all management stations you have set up for SSH access to the switch using the earlier pair.

Removing (zeroizing) the switch's public/private key pair renders the switch unable to engage in SSH operation and automatically disables IP SSH on the switch. (To verify whether SSH is enabled, execute **show ip ssh**.)

To Generate or Erase the Switch's Public/Private RSA Host Key Pair. Because the host key pair is stored in flash instead of the running-config file, it is not necessary to use **write memory** to save the key pair. Erasing the key pair automatically disables SSH.

Syntax:	<code>crypto key generate [rsa]</code>	<i>Generates a public/private key pair for the switch. If a switch key pair already exists, replaces it with a new key pair. (See the Note, above.)</i>
	<code>crypto key zeroize [rsa]</code>	<i>Erases the switch's public/private key pair and disables SSH operation.</i>
	<code>show ip ssh host-public-key [babble] [fingerprint]</code>	<i>Displays switch's public key as an ASCII string. Displays a hash of the switch's public key in phonetic format. (See "Displaying the Public Key" on page 73.) Displays a "fingerprint" of the switch's public key in hexadecimal format. (See "Displaying the Public Key" on page 73.)</i>

For example, to generate and display a new key:

```

HPswitch(config)# crypto key generate rsa
Generating new RSA host key.  If the cache is depleted,
this could take up to two minutes.
HPswitch(config)# show ip host-public-key
896 35 4271994707660774263666250605799242148515279332487520218551264932934075407
04782860432930458032140273304999167004670769854352973485302001767770553555445568
80992231580238056056245444224389955500310200336191361046978602009243623264937429
4060627777506601747146563337525446401

```

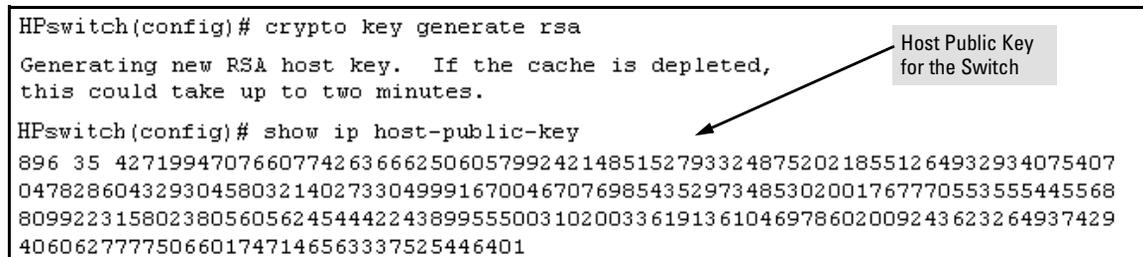


Figure 25. Example of Generating a Public/Private Host Key Pair for the Switch

Notes

"Zeroizing" the switch's key automatically disables SSH (sets **IP SSH** to **No**). Thus, if you zeroize the key and then generate a new key, you must also re-enable SSH with the **ip ssh** command before the switch can resume SSH operation.

3. Providing the Switch's Public Key to Clients

When an SSH client contacts the switch for the first time, the client will challenge the connection unless you have already copied the key into the client's "known host" file. Copying the switch's key in this way reduces the chance that an unauthorized device can pose as the switch to learn your access passwords. The most secure way to acquire the switch's public key for distribution to clients is to use a direct, serial connection between the switch and a management device (laptop, PC, or UNIX workstation), as described below.

Note on the Public Key Format

The switch uses SSH version 1, but can be authenticated by SSH version 2 clients that are backwards-compatible to SSHv1. However, if your SSH client supports SSHv2, then it may use the PEM format for storing the switch's public key in its "known host" file. In this case, the following procedure will not work for the client unless you have a method for converting the switch's ASCII-string public key into the PEM format. If you do not have a conversion method, then you can still set up authentication of the switch to the client over the network by simply using your client to contact the switch and then accepting the resulting challenge that your client should pose to accepting the switch. This should be acceptable as long as you are confident that there is no "man-in-the-middle" spoofing attempt during the first contact. Because the client will acquire the switch's public key after you accept the challenge, subsequent contacts between the client and the switch should be secure.

The public key generated by the switch consists of three parts, separated by one blank space each:

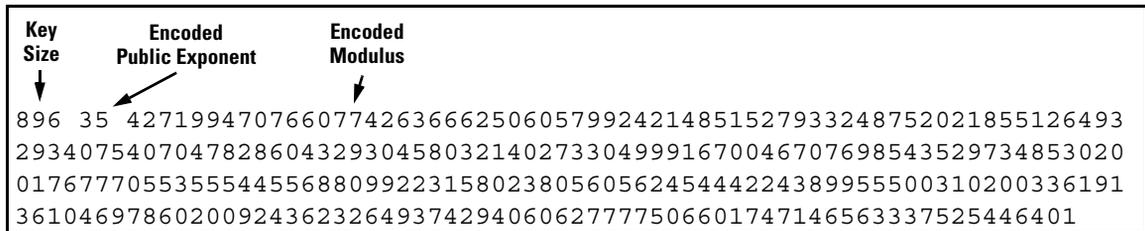


Figure 26. Example of a Public Key Generated by the Switch

(The generated public key on the switch is always 896 bits.)

With a direct serial connection from a management station to the switch:

1. Use a terminal application such as HyperTerminal to display the switch's public key with the **show ip host-public-key** command, as shown in figure 25.
2. Bring up the SSH client's "known host" file in a text editor such as Notepad as straight ASCII text, and copy the switch's public key into the file.

- Ensure that there are no line breaks in the text string. (A public key must be an unbroken ASCII string. Line breaks are not allowed.) For example, if you are using Windows® Notepad, ensure that **Word Wrap** (in the **Edit** menu) is disabled, and that the key text appears on a single line.

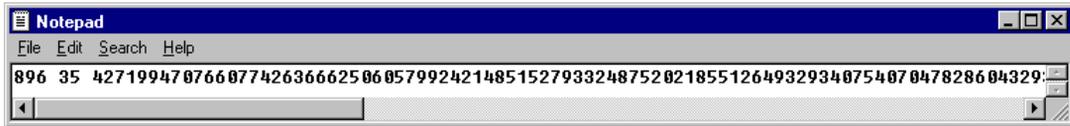


Figure 27. Example of a Correctly Formatted Public Key (Unbroken ASCII String)

- Add any data required by your SSH client application. For example Before saving the key to an SSH client's "known hosts" file you may have to insert the switch's IP address:

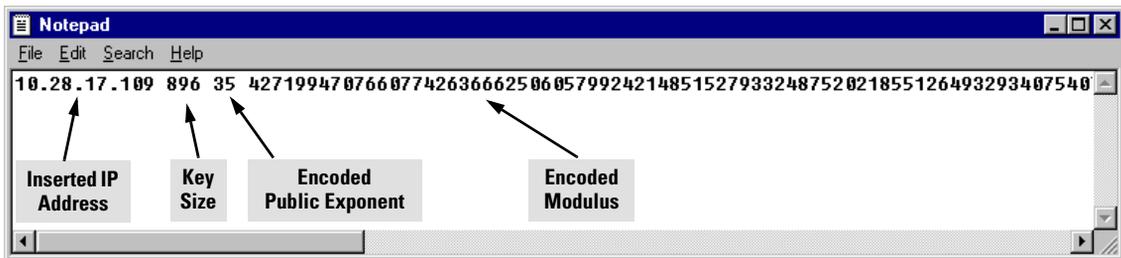


Figure 28. Example of a Switch Public Key Edited To Include the Switch's IP Address

For more on this topic, refer to the documentation provided with your SSH client application.

Displaying the Public Key. The switch provides three options for displaying its public key. This is helpful if you need to visually verify that the public key the switch is using for authenticating itself to a client matches the copy of this key in the client's "known hosts" file:

- **Non-encoded ASCII numeric string:** Requires a client ability to display the keys in the "known hosts" file in the ASCII format. This method is tedious and error-prone due to the large ASCII number set. (See figure 27 on page 73.)
- **Phonetic hash:** Outputs the key as a relatively short series of alphabetic character groups. Requires a client ability to convert the key to this format.
- **Hexadecimal hash:** Outputs the key as a relatively short series of hexadecimal numbers. Requires a parallel client ability.

For example, on the switch, you would generate the phonetic and hexadecimal versions of the switch's public key in figure 27 as follows:

```
HP2512# show ip host-public-key babble
896 xurac-metym-sagaf-recus-caleb-niten-kames-pobud-poluc-gelyl-exxas

HP2512# show ip host-public-key fingerprint
896 f3:f3:61:4c:06:ea:53:c2:7f:e7:78:b6:9b:7f:88:5b
```

Phonetic "Hash" of Switch's Public Key

Hexadecimal "Hash" of the Same Switch Public Key

Figure 29. Examples of Visual Phonetic and Hexadecimal Conversions of the Switch's Public Key

Note

The two commands shown in figure 29 convert the displayed format of the switch's (host) public key for easier visual comparison of the switch's public key to a copy of the key in a client's "known host" file. The switch always uses an ASCII version (without PEM encoding, or babble or fingerprint conversion) of its public key for file storage and default display format.

4. Enabling SSH on the Switch and Anticipating SSH Client Contact Behavior

The **ip ssh** command enables or disables SSH on the switch and modifies parameters the switch uses for transactions with clients. After you enable SSH, the switch can authenticate itself to SSH clients.

Note

Before enabling SSH on the switch you must generate the switch's public/private key pair. If you have not already done so, refer to "2. Generating the Switch's Public and Private Key Pair" on page 70.

When configured for SSH, the switch uses its host public-key to authenticate itself to SSH clients. If you also want SSH clients to authenticate themselves to the switch you must do one of the following:

- Configure SSH on the switch for client public-key authentication at the login (Operator) level, with (optionally) local, TACACS+, or RADIUS authentication at the enable (Manager) level.
- Configure SSH on the switch for local, TACACS+, or RADIUS password authentication at the login and enable levels.

Refer to "5. Configuring the Switch for SSH Authentication" on page 77.

SSH Client Contact Behavior. At the first contact between the switch and an SSH client, if you have not copied the switch's public key into the switch, your client's first connection to the switch will question the connection and, for security reasons, give you the option of accepting or refusing. As long as you are confident that an unauthorized device is not using the switch's IP address in an attempt to gain access to your data or network, you can accept the connection. (As a more secure alternative, you can directly connect the client to the switch's serial port and copy the switch's public key into the client. See the Note, below.)

Note

When an SSH client connects to the switch for the first time, it is possible for a "man-in-the-middle" attack; that is, for an unauthorized device to pose undetected as the switch, and learn the usernames and passwords controlling access to the switch. You can remove this possibility by directly connecting the management station to the switch's serial port, using a **show** command to display the switch's public key, and copying the key from the display into a file. This requires a knowledge of where your client stores public keys, plus the knowledge of what key editing and file format might be required by your client application. However, if your first contact attempt between a client and the switch does not pose a security problem, this is unnecessary.

To enable SSH on the switch.

1. Generate a public/private key pair if you have not already done so. (Refer to "2. Generating the Switch's Public and Private Key Pair" on page 70.)
2. Execute the **ip ssh** command.

To disable SSH on the switch, do either of the following:

- Execute **no ip ssh**.
- Zeroize the switch's existing key pair. (page 71).

Syntax: [no] ip ssh	<i>Enables or disables SSH on the switch.</i>
[key-size < 512 768 1024 >]	<i>The size of the internal, automatically generated key the switch uses for negotiations with an SSH client. A larger key provides greater security; a smaller key results in faster authentication (default: 512 bits). See the following Note.</i>
[port < 1-65535 default >]	<i>The IP port number for SSH connections (default: 22). Important: See the following "Note" on port number.</i>
[timeout < 5 - 120 >]	<i>The SSH login timeout value (default: 120 seconds).</i>

Note on Port Number

The **ip ssh key-size** command affects only a per-session, internal server key the switch creates, uses, and discards. This key is not accessible from the user interface. The switch's public (host) key is a separate, accessible key that is always 896 bits.

HP recommends using the default IP port number (22). However, you can use **ip ssh port** to specify any TCP port for SSH connections except those reserved for other purposes. Examples of reserved IP ports are 23 (Telnet) and 80 (http). Some other commonly reserved IP ports are 49, 80, 1506, and 1513.

```

HPswitch(config)# ip ssh
HPswitch(config)# show ip ssh

SSH Enabled           : Yes

IP Port Number       : 22
Timeout (sec)       : 120
Server Key Size (bits) : 512

Ses Type      Source IP and Port
-----
1  console
2  ssh        15.30.252.195:1722
3  telnet
4  inactive
    
```

← Enables SSH on the switch.

← Lists the current SSH configuration and status.

← The switch uses these three settings internally for transactions with clients. See the **Note**, below.

With SSH running, the switch allows one console session and up to three other sessions (SSH and/or Telnet). Web browser sessions are also allowed, but does not appear in the **show ip ssh** listing.

Figure 30. Example of Enabling IP SSH and Listing the SSH Configuration and Status

Caution

Protect your private key file from access by anyone other than yourself. If someone can access your private key file, they can then penetrate SSH security on the switch by appearing to be you.

SSH does not protect the switch from unauthorized access via the web interface, Telnet, SNMP, or the serial port. While web and Telnet access can be restricted by the use of passwords local to the switch, if you are unsure of the security this provides, you may want to disable web-based and/or Telnet access (**no web-management** and **no telnet**). If you need to increase SNMP security, use the **snmp security** command. Another security measure is to use the Authorized IP Managers feature described in the switch's *Management and Configuration Guide*. To protect against unauthorized access to the serial port (and the Clear button, which removes local password protection), keep physical access to the switch restricted to authorized personnel.


```

HP2512(config)# password manager user-name leader
New password for Manager: *****
Please retype new password for Manager: *****

HP2512(config)# aaa authentication ssh login rsa none
HP2512(config)# aaa authentication ssh enable tacacs local
HP2512(config)# copy tftp pub-key-file 10.33.18.117 Client-Keys.pub
HP2512(config)# write memory
    
```

Configures Manager user-name and password.

Configures the switch to allow SSH access only a client whose public key matches one of the keys in the public key file downloaded to the switch.

Copies a public key file named "Client-Keys.pub" into the switch.

Configures the primary and secondary password methods for Manager (Becomes available after SSH access is granted to a client.)

Figure 31. Configuring for SSH Access Requiring a Client Public-Key Match and Manager Passwords

Figure 32 shows how to check the results of the above commands.

```

HP2512(config)# show authentication
Status and Counters - Authentication Information
Login Attempts : 3

  Access Task | Login      Login      Enable     Enable
              | Primary   Secondary  Primary    Secondary
  -----+-----+-----+-----+-----
  Console     | Local     None       Local       None
  Telnet      | Local     None       Local       None
  Port-Access | Local
  SSH         | RSA       None       Tacacs      Local
    
```

Lists the current SSH authentication configuration.

```

HP2512(config)# show ip client-public-key
1024 35
1140740666170144690796380365284018053912704374511148288250978555011016860308261603
9146896306569035982041222025542543282764329943344032963504781021098947647460524651
6455722276820316076486036640205347034083710028842932315034922654093553211199274541
543765609589968291386053556814705585051025488575846923smith@fellow

512 35
9210552662725956416192357815130947052205817856719413127292584858843442080864040396
015631513328914504076264918047285352382655581475615051610646634228991kjlwilson@gray
lds
    
```

Shows the contents of the public key file downloaded with the copy tftp command in figure 31. In this example, the file contains two client public-keys.

Figure 32. SSH Configuration and Client-Public-Key Listing From Figure 31

6. Use an SSH Client To Access the Switch

Test the SSH configuration on the switch to ensure that you have achieved the level of SSH operation you want for the switch. If you have problems, refer to “Troubleshooting SSH Operation” on page 86 for possible solutions.

Further Information on SSH Client Public-Key Authentication

The section titled “5. Configuring the Switch for SSH Authentication” on page 77 lists the steps for configuring SSH authentication on the switch. However, if you are new to SSH or need more details on client public-key authentication, this section may be helpful.

When configured for SSH operation, the switch automatically attempts to use its own host public-key to authenticate itself to SSH clients. To provide the optional, opposite service—client public-key authentication to the switch—you can configure the switch to authenticate up to ten SSH clients. This requires storing an ASCII version of each client’s public key (without PEM encoding, babble conversion, or fingerprint conversion) in a client public-key file that you create and TFTP-copy to the switch. In this case, only clients that have a private key corresponding to one of the stored public keys can gain access to the switch using SSH. *That is, if you use this feature, only the clients whose public keys are in the client public-key file you store on the switch will have SSH access to the switch over the network.* If you do not allow secondary SSH login (Operator) access via local password, then the switch will refuse other SSH clients.

SSH clients that support client public-key authentication normally provide a utility to generate a key pair. The private key is usually stored in a password-protected file on the local host; the public key is stored in another file and is not protected.

(Note that even without using client public-key authentication, you can still require authentication from whoever attempts to access the switch from an SSH client— by employing the local username/password, TACACS+, or RADIUS features. Refer to “5. Configuring the Switch for SSH Authentication” on page 77.)

If you enable client public-key authentication, the following events occur when a client tries to access the switch using SSH:

1. The client sends its public key to the switch with a request for authentication.
2. The switch compares the client’s public key to those stored in the switch’s client-public-key file. (As a prerequisite, you must use the switch’s **copy tftp** command to download this file to flash.)
3. If there is not a match, and you have not configured the switch to accept a login password as a secondary authentication method, the switch denies SSH access to the client.
4. If there is a match, the switch:
 - a. Generates a random sequence of bytes.
 - b. Uses the client’s public key to encrypt this sequence.
 - c. Send these encrypted bytes to the client.
5. The client uses its private key to decrypt the byte sequence.
6. The client then:
 - a. Combines the decrypted byte sequence with specific session data.

- b. Uses MD5 to create a hash version of this information.
 - c. Returns the hash version to the switch.
7. The switch computes its own hash version of the data in step 6 and compares it to the client's hash version. If they match, then the client is authenticated. Otherwise, the client is denied access.

Using client public-key authentication requires these steps:

1. Generate a public/private key pair for each client you want to have SSH access to the switch. This can be a separate key for each client or the same key copied to several clients.
2. Copy the public key for each client into a client-public-key text file. (For the SSHv1 application used in the switch, this must be in the ASCII format (without PEM or any other encoding). If you are using an SSHv2 client application that creates its public key in a PEM-encoded ASCII string, you will need to convert the client's public key to a non-encoded version. Refer to the documentation provided with the application.)
3. Use **copy tftp** to copy the client-public-key file into the switch. Note that the switch can hold only one of these files. If there is already a client-public-key file in the switch and you copy another one into the switch, the second file replaces the first file.
4. Use the **aaa authentication ssh** command to enable client public-key authentication.

To Create a Client-Public-Key Text File. These steps describe how to copy client-public-keys into the switch for RSA challenge-response authentication, and require an understanding of how to use your SSH client application.

Bit Size	Public Index	Modulus	Comment
1024	35	11407406661701446907963803652840180539127043745111482882509285550110168603082603895914689630656903598204122202554254328276432994334403296350438102109894764746056455722276820316076486036640205347034083710028842932315034982165409355321119922465153140745413543765609589968291386053556814705585051025488575846923	smith@support.cairns.com

Figure 33. Example of a Client Public Key

Notes

Comments in public key files, such as `smith@support.cairns.com` in figure 33, may appear in a SSH client application's generated public key. While such comments may help to distinguish one key from another, they do not pose any restriction on the use of a key by multiple clients and/or users.

Public key illustrations such as the key shown in figure 33 usually include line breaks as a method for showing the whole key. However, in practice, line breaks in a public key will cause errors resulting in authentication failure.

1. Use your SSH client application to create a public/private key pair. Refer to the documentation provided with your SSH client application for details. The Series 2500 switches support the following client-public-key properties:

Property	Supported Value	Comments
Key Format	ASCII (no PEM or other encoding)	See figure 27 on page 73. The key must be one unbroken, non-encoded ASCII string. If you add more than one client-public-key to a file, terminate each key (except the last one) with a <CR><LF>. Spaces are allowed within the key to delimit the key's components. Also, the switch supports only SSH version 1. If your SSH client supports SSHv2, then it may use the PEM format for creating its public key. In this case, you will need a method for converting the switch's PEM-formatted public key into an ASCII-string equivalent. Note that, unlike the use of the switch's public key in an SSH client application, the format of a client-public-key used by the switch does not include the client's IP address.
Key Type	RSA only	
Maximum Supported Public Key Length	3072 bits	Shorter key lengths allow faster operation, but also mean diminished security.
Maximum Key Size	1024 characters	Includes the bit size, public index, modulus, any comments, <CR>, <LF>, and all blank spaces. If necessary, you can use an editor application to verify the size of a key. For example, if you place a client-public-key into a Word for Windows text file and then click on File Properties Statistics , you can view the number of characters in the file, including spaces.

2. Copy the client's public key (in ASCII, non-encoded format) into a text file (*filename.txt*). (For example, you can use the Notepad editor included with the Microsoft® Windows® software. If you want several clients to use client public-key authentication, copy a public key for each of these clients (up to ten) into the file. Each key should be separated from the preceding key by a <CR><LF>.
3. Copy the client-public-key file into a TFTP server accessible to the switch.

Copying a client-public-key into the switch requires the following:

- One or more client-generated public keys in non-encoded ASCII format. If you are using an SSHv2 client application, a client may encode its public key in PEM format. *To use the client public-key feature, you will need to convert the key to a non-encoded ASCII format.* Refer to the documentation provided with your SSH client application.
- A copy of each client public key (up to ten) stored in a single text file on a TFTP server to which the switch has access. (The text file should contain all client public keys for the clients you want to have access to the switch.) Terminate all client public-keys in the file except the last one with a <CR><LF>.

Note on Public Keys

The actual content of a public key entry in a public key file is determined by the SSH client application generating the key. (Although you can manually add or edit any comments the client application adds to the end of the key, such as the `smith@fellow` at the end of the key in figure 33, above.)

The file on the TFTP server must contain non-encoded ASCII text of each public key you want copied. Also, the file must be a text file (such as `filename.txt`).

Syntax: `copy tftp pub-key-file <ip-address> <filename>`

Copies a public key file from a TFTP server into flash memory in the switch.

`show ip client-public-key [babble | fingerprint]`

Displays the client public key(s) in the switch's current client-public-key file.

*The **babble** option converts the key data to a phonetic hash that is easier for visual comparisons.*

*The **fingerprint** option converts the key data to a hexadecimal hash for the same purpose.*

For example, if you wanted to copy a client public-key file named `clientkeys.txt` from a TFTP server at `10.38.252.195` and then display the file contents:

```
HP2512(config)# copy tftp pub-key-file 10.38.252.195 clientkeys.txt
HP2512(config)# show ip client-public-key
1024 35
1140740666170144690796380365284018053912704374511148288250928555011016860308261603895
9146896306569035982041222025542543282764329943344032963504381021098947647460524651531
6455722276820316076486036640205347034083710028842932315034922654093553211199274541340
543765609589968291386053556814705585051025488575846923smith@fellow

512 35
9210552662725956416192357815130947052205817856719413127292584858843442080864040396207
015631513328914504076264918047285352382655581475615051610646634228991kjwilson@grayfile
lds
```

Figure 34. Example of Copying and Displaying a Client Public-Key File Containing Two Client Public Keys

Replacing or Clearing the Public Key File. The client public-key file remains in the switch's flash memory even if you erase the startup-config file, reset the switch, or reboot the switch.

- You can replace the existing client public-key file by copying a new client public-key file into the switch
- You can remove the existing client public-key file by executing the **clear public-key** command.

Syntax: clear public-key *Deletes the client-public-key from the switch.*

For example:

```
HP2512(config)# clear public-key
HP2512(config)# show ip client-public-key
show_client_public_key: cannot stat keyfile
```

Clearing the public key file removes file from flash memory, and does not require a write memory command to make the change permanent.

Enabling Client Public-Key Authentication. After you TFTP a client-public-key file into the switch (described above), you can configure the switch to allow one of the following:

- If an SSH client's public key matches the switch's client-public-key file, allow that client access to the switch. If there is not a public-key match, then deny access to that client.
- If an SSH client's public key does not have a match in the switch's client-public-key file, allow the client access if the user can enter the switch's login (Operator) password. (If the switch does not have an Operator password, then deny access to that client.

Syntax: aaa authentication ssh login rsa none *Allows SSH client access only if the switch detects a match between the client's public key and an entry in the client-public-key file most recently copied into the switch.*

 aaa authentication ssh login rsa local *Allows SSH client access if there is a public key match (see above) or if the client's user enters the switch's login (Operator) password.*

With **login rsa local** configured, if the switch does not have an Operator-level password, it blocks client public-key access to SSH clients whose private keys do not match a public key in the switch's client-public-key file.

Caution

To enable client public-key authentication to block SSH clients whose public keys are not in the client-public-key file copied into the switch, you must configure the Login Secondary as **none**. Otherwise, the switch allows such clients to attempt access using the switch's Operator password.

Messages Related to SSH Operation

Message	Meaning
00000K Peer unreachable.	<p>Indicates an error in communicating with the tftp server or not finding the file to download. Causes include such factors as:</p> <ul style="list-style-type: none"> • Incorrect IP configuration on the switch • Incorrect IP address in the command • Case (upper/lower) error in the filename used in the command • Incorrect configuration on the TFTP server • The file is not in the expected location. • Network misconfiguration • No cable connection to the network
00000K Transport error.	<p>Indicates the switch experienced a problem when trying to copy tftp the requested file. The file may not be in the expected directory, the filename may be misspelled in the command, or the file permissions may be wrong.</p>
Cannot bind reserved TCP port <port-number>.	<p>The ip ssh port command has attempted to configure a reserved TCP port. Use the default or select another port number. See "Note on Port Number" on page 76.</p>
Client public key file corrupt or not found. Use 'copy tftp pub-key-file <ip-addr> <filename>' to download new file.	<p>The client key does not exist in the switch. Use copy tftp to download the key from a TFTP server.</p>
Download failed: overlength key in key file.	<p>The public key file you are trying to download has one of the following problems:</p> <ul style="list-style-type: none"> • 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. <p>Refer to "To Create a Client-Public-Key Text File" on page 81 for information on client-public-key properties.</p>
Download failed: too many keys in key file.	
Download failed: one or more keys is not a valid RSA public key.	
Error: Requested keyfile does not exist.	<p>The client key does not exist in the switch. Use copy tftp to download the key from a TFTP server.</p>

Message	Meaning
Generating new RSA host key. If the cache is depleted, this could take up to two minutes.	After you execute the <code>crypto key generate [rsa]</code> command, the switch displays this message while it is generating the key.
Host RSA key file corrupt or not found. Use 'crypto key generate rsa' to create new host key.	The switch's key is missing or corrupt. Use the crypto key generate [rsa] command to generate a new key for the switch.
host_ssh1 is not a valid key file. Key does not exist or is corrupt. show client_public-key: cannot stat keyfile.	The client key does not exist in the switch. Use copy tftp to download the key from a TFTP server.

Troubleshooting SSH Operation

See also “Messages Related to SSH Operation” on page 85.

Symptom	Possible Cause
Switch access refused to a client whose public key you have placed in a text file and copied (using the copy tftp public-key-file command) into the switch.	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 <code>show ip host-public-key</code> . If you see the message ssh cannot be enabled until a host key is configured (use 'crypto' command) then you need to generate an SSH key pair for the switch. To do so, execute crypto key generate .(Refer to “2. Generating the Switch's Public and Private Key Pair” on page 70.)
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 newline (CR). While this is optional for the last entry in the file, not adding a newline 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.

Symptom	Possible Cause
<p>An attempt to copy a client public-key file into the switch has failed and the switch lists one of the following messages:</p> <p>Download failed: overlenght key in key file.</p> <p>Download failed: too many keys in key file.</p> <p>Download failed: one or more keys is not a valid RSA public key.</p>	<p>The public key file you are trying to download has one of the following problems:</p> <ul style="list-style-type: none"> • 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.
<p>Client ceases to respond ("hangs") during connection phase.</p>	<p>The switch does not support data compression in an SSH session. Clients will often have compression turned on by default, but will disable it during the negotiation phase. A client which 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.</p>

— This page is intentionally unused. —

Configuring RADIUS Authentication and Accounting

Feature	Default	Menu	CLI	Web
Configuring RADIUS Authentication	None	n/a	page 92	n/a
Configuring RADIUS Accounting	None	n/a	page 102	n/a
Viewing RADIUS Statistics	n/a	n/a	page 108	n/a

RADIUS (*Remote Authentication Dial-In User Service*) enables you to use up to three servers (one primary server and one or two backups) and maintain separate authentication and accounting for each RADIUS server employed. For authentication, this allows a different password for each user instead of having to rely on maintaining and distributing switch-specific passwords to all users. For accounting, this can help you track network resource usage.

Authentication. You can use RADIUS to verify user identity for the following types of primary password access to the Series 2500 switches:

- Serial port (Console)
- Telnet
- SSH
- Port-Access

Note

The Series 2500 switches do not support RADIUS security for SNMP (network management) access or web browser interface access. For steps to block unauthorized access through the web browser interface, see “Controlling Web Browser Interface Access When Using RADIUS Authentication” on page 101.

Accounting. RADIUS accounting on the Series 2500 switches collects resource consumption data and forwards it to the RADIUS server. This data can be used for trend analysis, capacity planning, billing, auditing, and cost analysis.

Terminology

CHAP (Challenge-Handshake Authentication Protocol): A challenge-response authentication protocol that uses the Message Digest 5 (MD5) hashing scheme to encrypt a response to a challenge from a RADIUS server.

EAP(Extensible Authentication Protocol): A general PPP authentication protocol that supports multiple authentication mechanisms. A specific authentication mechanism is known as an EAP type, such as MD5-Challenge, Generic Token Card, and TLS (Transport Level Security).

Host: See **RADIUS Server**.

NAS (Network Access Server): In this case, a Switch 2512 or 2524 configured for RADIUS security operation.

RADIUS (Remote Authentication Dial In User Service):

RADIUS Client: The device that passes user information to designated RADIUS servers.

RADIUS Host: See RADIUS server.

RADIUS Server: A server running the RADIUS application you are using on your network. This server receives user connection requests from the switch, authenticates users, and then returns all necessary information to the switch. For the Switch 2512 and 2524 a RADIUS server can also perform accounting functions. Sometimes termed a **RADIUS host**.

Shared Secret Key: A text value used for encrypting data in RADIUS packets. Both the RADIUS client and the RADIUS server have a copy of the key, and the key is never transmitted across the network.

Switch Operating Rules for RADIUS

- You must have at least one RADIUS server accessible to the switch.
- The switch supports authentication and accounting using up to three RADIUS servers. The switch accesses the servers in the order in which they are listed by the **show radius** command (page 108). If the first server does not respond, the switch tries the next one, and so-on. (To change the order in which the switch accesses RADIUS servers, refer to “Changing RADIUS-Server Access Order” on page 113.)
- You can select RADIUS as the primary authentication method for each type of access. (Only one primary and one secondary access method is allowed for each access type.)
- In the Series 2500 switches, EAP RADIUS uses MD5 and TLS to encrypt a response to a challenge from a RADIUS server.

General RADIUS Setup Procedure

Preparation:

1. Configure one to three RADIUS servers to support the switch. (That is, one primary server and one or two backups.) Refer to the documentation provided with the RADIUS server application.
2. Before beginning to configure the switch, collect the information outlined below.

Table 7. Preparation for Configuring RADIUS on the Switch

- Determine the access methods (console, Telnet, Port-Access, and/or SSH) for which you want RADIUS as the primary authentication method. Consider both Operator (login) and Manager (enable) levels, as well as which secondary authentication methods to use (local or none) if the RADIUS authentication fails or does not respond.

```
HP2512 (config)# show authentication
Status and Counters - Authentication Information
Login Attempts : 3
```

Access Task	Login Primary	Login Secondary	Enable Primary	Enable Secondary
Console	Radius	Local	Radius	Local
Telnet	Radius	None	Radius	None
Port-Access	EapRadius			
SSH	Radius	None	Radius	None

Console access requires Local as secondary method to prevent lockout if the primary RADIUS access fails due to loss of RADIUS server access or other problems with the server.

Figure 35. Example of Possible RADIUS Access Assignments

- Determine the IP address(es) of the RADIUS server(s) you want to support the switch. (You can configure the switch for up to three RADIUS servers.)
- If you need to replace the default UDP destination port (1812) the switch uses for authentication requests to a specific RADIUS server, select it before beginning the configuration process.
- If you need to replace the default UDP destination port (1813) the switch uses for accounting requests to a specific Radius server, select it before beginning the configuration process.
- Determine whether you can use one, global encryption key for all RADIUS servers or if unique keys will be required for specific servers. With multiple RADIUS servers, if one key applies to two or more of these servers, then you can configure this key as the global encryption key. For any server whose key differs from the global key you are using, you must configure that key in the same command that you use to designate that server's IP address to the switch.
- Determine an acceptable timeout period for the switch to wait for a server to respond to a request. HP recommends that you begin with the default (five seconds).
- Determine how many times you want the switch to try contacting a RADIUS server before trying another RADIUS server or quitting. (This depends on how many RADIUS servers you have configured the switch to access.)
- Determine whether you want to bypass a RADIUS server that fails to respond to requests for service. To shorten authentication time, you can set a bypass period in the range of 1 to 1440 minutes for non-responsive servers. This requires that you have multiple RADIUS servers accessible for service requests.

Configuring the Switch for RADIUS Authentication

RADIUS Authentication Commands	
aaa authentication	page 94
< console telnet ssh > < enable login > radius	page 94
< local none >	page 94
[no] radius-server host < <i>IP-address</i> >	page 96
[auth-port < <i>port-number</i> >]	page 96
[acct-port < <i>port-number</i> >]	page 96, 105
[key < <i>server-specific key-string</i> >]	page 96
[no] radius-server key < <i>global key-string</i> >	page 98
radius-server timeout < 1 .. 15 >	page 98
radius-server retransmit < 1 .. 5 >	page 98
[no] radius-server dead-time < 1 .. 1440 >	page 99
show radius	page 108
[< host < <i>ip-address</i> >]	page 109
show authentication	page 111
show radius authentication	page 111

Outline of the Steps for Configuring RADIUS Authentication

There are three main steps to configuring RADIUS authentication:

1. Configure RADIUS authentication for controlling access through one or more of the following
 - Serial port
 - Telnet
 - SSH
 - Port-Access (802.1x)
2. Configure the switch for accessing one or more RADIUS servers (one primary server and up to two backup servers):

Note

This step assumes you have already configured the RADIUS server(s) to support the switch. Refer to the documentation provided with the RADIUS server documentation.)

- Server IP address
 - (Optional) UDP destination port for authentication requests (default: 1812; recommended)
 - (Optional) UDP destination port for accounting requests (default: 1813; recommended)
 - (Optional) encryption key for use during authentication sessions with a RADIUS server. This key overrides the global encryption key you can also configure on the switch, and must match the encryption key used on the specified RADIUS server. (Default: null)
3. Configure the global RADIUS parameters.
 - **Server Key:** This key must match the encryption key used on the RADIUS servers the switch contacts for authentication and accounting services unless you configure one or more per-server keys. (Default: null.)
 - **Timeout Period:** The timeout period the switch waits for a RADIUS server to reply. (Default: 5 seconds; range: 1 to 15 seconds.)
 - **Retransmit Attempts:** The number of retries when there is no server response to a RADIUS authentication request. (Default: 3; range of 1 to 5.)
 - **Server Dead-Time:** The period during which the switch will not send new authentication requests to a RADIUS server that has failed to respond to a previous request. This avoids a wait for a request to time out on a server that is unavailable. If you want to use this feature, select a dead-time period of 1 to 1440 minutes. (Default: 0—disabled; range: 1 - 1440 minutes.) If your first-choice server was initially unavailable, but then becomes available before the dead-time expires, you can nullify the dead-time by resetting it to

zero and then trying to log on again. As an alternative, you can reboot the switch, (thus resetting the dead-time counter to assume the server is available) and then try to log on again.

- **Number of Login Attempts:** This is actually an **aaa authentication** command. It controls how many times in one session a RADIUS client (as well as clients using other forms of access) can try to log in with the correct username and password. (Default: Three times per session.)

(For RADIUS accounting features, refer to “Configuring RADIUS Accounting” on page 102.)

1. Configure Authentication for the Access Methods You Want RADIUS To Protect

This section describes how to configure the switch for RADIUS authentication through the following access methods:

- **Console:** Either direct serial-port connection or modem connection.
- **Telnet:** Inbound Telnet must be enabled (the default).
- **SSH:** To employ RADIUS for SSH access, you must first configure the switch for SSH operation. Refer to “Configuring Secure Shell (SSH)” on page 63.

You can also use RADIUS for Port-Based Access authentication. Refer to “Configuring Port-Based Access Control (802.1x)” on page 65.

You can configure RADIUS as the primary password authentication method for the above access methods. You will also need to select either **local** or **none** as a secondary, or backup, method. Note that for console access, if you configure **radius** (or **tacacs**) for primary authentication, you must configure **local** for the secondary method. This prevents the possibility of being completely locked out of the switch in the event that all primary access methods fail.

Syntax: aaa authentication < console | telnet | ssh >
 < enable | login > < radius >

[< local | none >]

*Configures RADIUS as the primary password authentication method for console, Telnet, and/or SSH. (The default primary < enable | login > authentication is **local**.)*

*Options for secondary authentication (default: **none**). Note that for console access, secondary authentication must be **local** if primary access is not **local**. This prevents you from being completely locked out of the switch in the event of a failure in other access methods.*

For example, suppose you have already configured local passwords on the switch, but want to use RADIUS to protect primary Telnet and SSH access without allowing a secondary Telnet or SSH access option (which would be the switch's local passwords):

```

HP2512(config)# aaa authentication telnet login radius none
HP2512(config)# aaa authentication telnet enable radius none
HP2512(config)# aaa authentication ssh login radius none
HP2512(config)# aaa authentication ssh enable radius none

HP2512(config)# show authentication
Status and Counters - Authentication Information
Login Attempts : 3

```

Access Task	Login Primary	Login Secondary	Enable Primary	Enable Secondary
Console	Local	None	Local	None
Telnet	Radius	None	Radius	None
Port-Access	Local			
SSH	Radius	None	Radius	None

The switch now allows Telnet and SSH authentication only through RADIUS.

Figure 36. Example Configuration for RADIUS Authentication

Note

In the above example, if you configure the Login Primary method as **local** instead of **radius** (and local passwords are configured on the switch), then you can gain access to either the Operator or Manager level without encountering the RADIUS authentication specified for Enable Primary. Refer to “Local Authentication Process” on page 100.

2. Configure the Switch To Access a RADIUS Server

This section describes how to configure the switch to interact with a RADIUS server for both authentication and accounting services. (If you want to configure RADIUS accounting on the switch, go to “Configuring RADIUS Accounting” on page 102 instead of continuing here.)

Syntax: [no] radius-server host < ip-address >	<i>Adds a server to the RADIUS configuration or (with no) deletes a server from the configuration. You can configure up to three RADIUS server addresses. The switch uses the first server it successfully accesses. (Refer to "Changing the RADIUS Server Access Order" on page 113.)</i>
[auth-port < port-number >]	<i>Optional. Changes the UDP destination port for authentication requests to the specified RADIUS server (host). If you do not use this option with the radius-server host command, the switch automatically assigns the default authentication port number. The auth-port number must match its server counterpart. (Default: 1812)</i>
[acct-port < port-number >]	<i>Optional. Changes the UDP destination port for accounting requests to the specified RADIUS server. If you do not use this option with the radius-server host command, the switch automatically assigns the default accounting port number. The acct-port number must match its server counterpart. (Default: 1813)</i>
[key < key-string >]	<i>Optional. Specifies an encryption key for use during authentication (or accounting) sessions with the specified server. This key must match the encryption key used on the RADIUS server. Use this command only if the specified server requires a different encryption key than configured for the global encryption key.</i>
no radius-server host < ip-address > key	<i>Use the no form of the command to remove the key for a specified server.</i>

For example, suppose you have configured the switch as shown in figure 37 and you now need to make the following changes:

1. Change the encryption key for the server at 10.33.18.127 to "source0127".
2. Add a RADIUS server with an IP address of 10.33.18.119 and a server-specific encryption key of "source0119".

```
HP2512(config)# show radius
Status and Counters - General RADIUS Information
Deadtme(min) : 0
Timeout(secs) : 5
Retransmit Attempts : 5
Global Encryption Key :
      Auth  Acct
Server IP Addr  Port  Port  Encryption Key
-----
10.33.18.127   1812 1813  TempKey01
```

Figure 37. Sample Configuration for RADIUS Server Before Changing the Key and Adding Another Server

To make the changes listed prior to figure 37, you would do the following:

```
HP2512(config)# radius-server host 10.33.18.127 key source0127
HP2512(config)# radius-server host 10.33.18.119 key source0119
HP2512(config)# show radius
Status and Counters - General RADIUS Information
Deadtme(min) : 0
Timeout(secs) : 5
Retransmit Attempts : 5
Global Encryption Key :
      Auth  Acct
Server IP Addr  Port  Port  Encryption Key
-----
10.33.18.127   1812 1813  source0127
10.33.18.119   1812 1813  source0119
```

Changes the key for the existing server to "source0127" (step 1, above).

Adds the new RADIUS server with its required "source0119" key.

Lists the switch's new RADIUS server configuration. Compare this with figure 37.

Figure 38. Sample Configuration for RADIUS Server After Changing the Key and Adding Another Server

To change the order in which the switch accesses RADIUS servers, refer to "Changing RADIUS-Server Access Order" on page 113.

3. Configure the Switch's Global RADIUS Parameters

You can configure the switch for the following global RADIUS parameters:

- **Number of login attempts:** In a given session, specifies how many tries at entering the correct username and password pair are allowed before access is denied and the session terminated. (This is a general **aaa authentication** parameter and is not specific to RADIUS.)
- **Global server key:** The server key the switch will use for contacts with all RADIUS servers for which there is not a server-specific key configured by **radius-server host** *< ip-address > key < key-string >*. This key is optional if you configure a server-specific key for each RADIUS server entered in the switch. (Refer to “2. Configure the Switch To Access a RADIUS Server” on page 96.)
- **Server timeout:** Defines the time period in seconds for authentication attempts. If the timeout period expires before a response is received, the attempt fails.
- **Server dead time:** Specifies the time in minutes during which the switch avoids requesting authentication from a server that has not responded to previous requests.
- **Retransmit attempts:** If the first attempt to contact a RADIUS server fails, specifies how many retries you want the switch to attempt on that server.

Syntax: aaa authentication num-attempts <1 .. 10 > *Specifies how many tries for entering the correct username and password before shutting down the session due to input errors. (Default: 3; Range: 1 - 10)*

[no] radius-server
 key < global-key-string > *Specifies the global encryption key the switch uses for sessions with servers for which the switch does not have a server-specific key assignment. This key is optional if all RADIUS server addresses configured in the switch include a server-specific encryption key. (Default: Null.)*

 dead-time < 1 .. 1440 > *Optional. Specifies the time in minutes during which the switch will not attempt to use a RADIUS server that has not responded to an earlier authentication attempt. (Default: 0; Range: 1 - 1440 minutes)*

radius-server timeout < 1 .. 15 > *Specifies the maximum time the switch waits for a response to an authentication request before counting the attempt as a failure. (Default: 3 seconds; Range: 1 - 15 seconds)*

radius-server retransmit < 1 .. 5 >

If a RADIUS server fails to respond to an authentication request, specifies how many retries to attempt before closing the session. (Default: 3; Range: 1 - 5)

Note

Where the switch has multiple RADIUS servers configured to support authentication requests, if the first server fails to respond, then the switch tries the next server in the list, and so-on. If none of the servers respond, then the switch attempts to use the secondary authentication method configured for the type of access being attempted (console, Telnet, or SSH). If this occurs, see “Troubleshooting RADIUS Operation” on page 115.

For example, suppose that your switch is configured to use three RADIUS servers for authenticating access through Telnet and SSH. Two of these servers use the same encryption key. In this case your plan is to configure the switch with the following global authentication parameters:

- Allow only two tries to correctly enter username and password.
- Use the global encryption key to support the two servers that use the same key. (For this example, assume that you did not configure these two servers with a server-specific key.)
- Use a dead-time of five minutes for a server that fails to respond to an authentication request.
- Allow three seconds for request timeouts.
- Allow two retries following a request that did not receive a response.

```
HP2512(config)# aaa authentication num-attempts 2
HP2512(config)# radius-server key My-Global-Key-1099
HP2512(config)# radius-server dead-time 5
HP2512(config)# radius-server timeout 3
HP2512(config)# radius-server retransmit 2
HP2512(config)# write mem
```

Figure 39. Example of Global Configuration Exercise for RADIUS Authentication

```

HP2512(config)# show authentication
Status and Counters - Authentication Information
( Login Attempts : 2 )
-----+-----
Access Task | Login      Login      Enable     Enable
              | Primary    Secondary  Primary    Secondary
-----+-----
Console     | Local      None       Local      None
Telnet      | Radius     Local      Radius     Local
Port-Access | Local
SSH         | Radius     Local      Radius     Local
    
```

After two attempts failing due to username or password entry errors, the switch will terminate the session.

```

HP2512(config)# show radius
Status and Counters - General RADIUS Information
( Deadttime(min) : 5
  Timeout(secs) : 3
  Retransmit Attempts : 2
  Global Encryption Key : My-Global-Key-1099 )
-----
Server IP Addr  Auth  Acct  Encryption Key
                Port  Port
-----
10.33.18.127    1812 1813  source0127
10.33.18.119    1812 1813
10.33.18.151    1812 1813
    
```

Global RADIUS parameters from figure 39.

Server-specific encryption key for the RADIUS server that will not use the global encryption key.

These two servers will use the global encryption key.

Figure 40. Listings of Global RADIUS Parameters Configured In Figure 39

Local Authentication Process

When the switch is configured to use RADIUS, it reverts to local authentication only if one of these two conditions exists:

- "Local" is the authentication option for the access method being used.
- The switch has been configured to query one or more RADIUS servers for a primary authentication request, but has not received a response, and local is the configured secondary option.

For local authentication, the switch uses the Operator-level and Manager-level username/password set(s) previously configured locally on the switch. (These are the usernames and passwords you can configure using the CLI password command, the web browser interface, or the menu interface—which enables only local password configuration).

- If the operator at the requesting terminal correctly enters the username/password pair for either access level (Operator or Manager), access is granted on the basis of which username/password pair was used. For example, suppose you configure Telnet primary access for RADIUS and Telnet secondary access for local. If a RADIUS access attempt fails, then you can still get access to either the Operator or Manager level of the switch by entering the correct username/password pair for the level you want to enter.
- If the username/password pair entered at the requesting terminal does not match either local username/password pair previously configured in the switch, access is denied. In this case, the terminal is again prompted to enter a username/password pair. In the default configuration, the switch allows up to three attempts. If the requesting terminal exhausts the attempt limit without a successful authentication, the login session is terminated and the operator at the requesting terminal must initiate a new session before trying again.

Controlling Web Browser Interface Access When Using RADIUS Authentication

Configuring the switch for RADIUS authentication does not affect web browser interface access. To prevent unauthorized access through the web browser interface, do one or more of the following:

- Configure local authentication (a Manager user name and password and, optionally, an Operator user name and password) on the switch.
- Configure the switch's Authorized IP Manager feature to allow web browser access only from authorized management stations. (The Authorized IP Manager feature does not interfere with TACACS+ operation.)
- Disable web browser access to the switch.

Configuring RADIUS Accounting

RADIUS Accounting Commands	
[no] radius-server host < ip-address >	page 105
[acct-port < port-number >]	page 105
[key < key-string >]	page 105
[no] aaa accounting < exec network system > < start-stop stop-only > radius	page 107
[no] aaa accounting update periodic < 1 .. 525600 > (in minutes)	page 108
[no] aaa accounting suppress null-username	page 108
show accounting	page 112
show accounting sessions	page 112
show radius accounting	page 112

Note

This section assumes you have already:

- Configured RADIUS authentication on the switch for one or more access methods
- Configured one or more RADIUS servers to support the switch

If you have not already done so, refer to “General RADIUS Setup Procedure” on page 91 before continuing here.

RADIUS accounting collects data about user activity and system events and sends it to a RADIUS server when specified events occur on the switch, such as a logoff or a reboot. The Series 2500 switches support three types of accounting services:

- **Network accounting:** Provides records containing the information listed below on clients directly connected to the switch and operating under Port-Based Access Control (802.1x):
 - Acct-Session-Id
 - Acct-Delay-Time
 - Nas-Port
 - Service-Type
 - Acct-Status-Type
 - Acct-Input-Packets
 - Acct-Output-Octets
 - NAS-IP-Address
 - Acct-Terminate-Cause
 - Acct-Output-Packets
 - Acct-Session-Time
 - NAS-Identifier
 - Acct-Authentic
 - Acct-Input-Octets
 - User-Name
 - Called-Station-Id

(For 802.1x information on the Series 2500 switches, refer to “Configuring Port-Based Access Control (802.1x)” on page 65.)

- **Exec accounting:** Provides records containing the information listed below about login sessions (console, Telnet, and SSH) on the switch:

- Acct-Session-Id
- Acct-Delay-Time
- NAS-IP-Address
- Acct-Status-Type
- Acct-Session-Time
- NAS-Identifier
- Acct-Terminate-Cause
- User-Name
- Calling-Station-Id
- Acct-Authentic
- Service-Type

- **System accounting:** Provides records containing the information listed below when system events occur on the switch, including system reset, system boot, and enabling or disabling of system accounting.

- Acct-Session-Id
- Acct-Delay-Time
- NAS-Identifier
- Acct-Status-Type
- User-Name
- Calling-Station-Id
- Acct-Terminate-Cause
- Service-Type
- Acct-Authentic
- NAS-IP-Address

The switch forwards the accounting information it collects to the designated RADIUS server, where the information is formatted, stored, and managed by the server. For more information on this aspect of RADIUS accounting, refer to the documentation provided with your RADIUS server.

Operating Rules for RADIUS Accounting

- You can configure up to three types of accounting to run simultaneously: exec, system, and network.
- RADIUS servers used for accounting are also used for authentication.
- The switch must be configured to access at least one RADIUS server.
- RADIUS servers are accessed in the order in which their IP addresses were configured in the switch. Use **show radius** to view the order. As long as the first server is accessible and responding to authentication requests from the switch, a second or third server will not be accessed. (For more on this topic, refer to “Changing RADIUS-Server Access Order” on page 113.)
- If access to a RADIUS server fails during a session, but after the client has been authenticated, the switch continues to assume the server is available to receive accounting data. Thus, if server access fails during a session, it will not receive accounting data transmitted from the switch.

Outline of the Steps for Configuring RADIUS Accounting

1. Configure the switch for accessing a RADIUS server.

You can configure a list of up to three RADIUS servers (one primary, two backup). The switch operates on the assumption that a server can operate in both accounting and authentication mode. (Refer to the documentation for your RADIUS server application.)

- Use the same **radius-server host** command that you would use to configure RADIUS authentication. Refer to “2. Configure the Switch To Access a RADIUS Server” on page 96.
- Provide the following:
 - A RADIUS server IP address.
 - Optional—a UDP destination port for authentication requests. Otherwise the switch assigns the default UDP port (1812; recommended).
 - Optional—if you are also configuring the switch for RADIUS authentication, and need a unique encryption key for use during authentication sessions with the RADIUS server you are designating, configure a server-specific key. This key overrides the global encryption key you can also configure on the switch, and must match the encryption key used on the specified RADIUS server. For more information, refer to the "[key < key-string >]" parameter on page 96. (Default: null)

2. Configure the types of accounting you want the switch to perform, and the controls for sending accounting reports from the switch to the RADIUS server(s).

- **Accounting types:** exec (page 103), network (page 102), or system (page 103)
- **Trigger for sending accounting reports to a RADIUS server:** At session start and stop or only at session stop

3. (Optional) Configure session blocking and interim updating options

- **Updating:** Periodically update the accounting data for sessions-in-progress
- **Suppress accounting:** Block the accounting session for any unknown user with no username accesses the switch

1. Configure the Switch To Access a RADIUS Server

Before you configure the actual accounting parameters, you should first configure the switch to use a RADIUS server. This is the same as the process described on page 96. You need to repeat this step here only if you have not yet configured the switch to use a RADIUS server, your server data has changed, or you need to specify a non-default UDP destination port for accounting requests. Note that switch operation expects a RADIUS server to accommodate both authentication and accounting.

Syntax: [no] radius-server host < ip-address >	<i>Adds a server to the RADIUS configuration or (with no) deletes a server from the configuration.</i>
[acct-port < port-number >]	<i>Optional. Changes the UDP destination port for accounting requests to the specified RADIUS server. If you do not use this option, the switch automatically assigns the default accounting port number. (Default: 1813)</i>
[key < key-string >]	<i>Optional. Specifies an encryption key for use during accounting or authentication sessions with the specified server. This key must match the encryption key used on the RADIUS server. Use this command only if the specified server requires a different encryption key than configured for the global encryption key.</i>

(For a more complete description of the **radius-server** command and its options, turn to page 96.)

For example, suppose you want the switch to use the RADIUS server described below for both authentication and accounting purposes.

- IP address: 10.33.18.151
- A non-default UDP port number of 1750 for accounting.

For this example, assume that all other RADIUS authentication parameters for accessing this server are acceptable at their default settings, and that RADIUS is already configured as an authentication method for one or more types of access to the switch (Telnet, Console, etc.).

```

HP2512(config)# radius-server host 10.33.18.151 acct-port 1750 key source0151
HP2512(config)# write mem
HP2512(config)# show radius

Status and Counters - General RADIUS Information
Deadtime(min) : 5
Timeout(secs) : 3
Retransmit Attempts : 2
Global Encryption Key :

Server IP Addr      Auth  Acct  Encryption Key
-----
10.33.18.151      1812 1750  source0151
    
```

Because the radius-server command includes an **acct-port** element with a non-default 1750, the switch assigns this value to the accounting port UDP port numbers. Because auth-port was not included in the command, the authentication UDP port is set to the default 1812.

Figure 41. Example of Configuring for a RADIUS Server with a Non-Default Accounting UDP Port Number

The radius-server command as shown in figure 41, above, configures the switch to use a RADIUS server at IP address 10.33.18.151, with a (non-default) UDP accounting port of 1750, and a server-specific key of "source0151".

2. Configure the Types of Accounting You Want the Switch to Perform, and the Controls for Sending Accounting Reports from the Switch to the RADIUS Server

Select the Accounting Type(s):

- **Exec:** Use **exec** if you want to collect accounting information on login sessions on the switch via the console, Telnet, or SSH. (See also “Accounting” on page 89.)
- **System:** Use **system** if you want to collect accounting data when:
 - A system boot or reload occurs
 - System accounting is turned on or off

Note that there is no timespan associated with using the **system** option. It simply causes the switch to transmit whatever accounting data it currently has when one of the above events occurs.

- **Network:** Use **Network** if you want to collect accounting information on 802.1x port-based-access users connected to the physical ports on the switch to access the network. (See also “Accounting” on page 89.) For information on this feature, refer to “Configuring Port-Based Access Control (802.1x)” on page 65.

Enhancements in Release F.04.08

- **Updates:** In addition to using a Start-Stop or Stop-Only trigger, you can optionally configure the switch to send periodic accounting record updates to a RADIUS server.
- **Suppress:** The switch can suppress accounting for an unknown user having no username.

Syntax: [no] aaa accounting update periodic < 1 .. 525600 > *Sets the accounting update period for all accounting sessions on the switch. (The **no** form disables the update function and resets the value to zero.) (Default: zero; disabled)*

[no] aaa accounting suppress null-username *Disables accounting for unknown users having no username. (Default: suppression disabled)*

To continue the example in figure 42, suppose that you wanted the switch to:

- Send updates every 10 minutes on in-progress accounting sessions.
- Block accounting for unknown users (no username).

```
HP2512(config)# aaa accounting update periodic 10
HP2512(config)# aaa accounting suppress null-username

HP2512(config)# show accounting
Status and Counters - Accounting Information
Interval(min) : 10
Suppress Empty User : Yes

Type   | Method Mode
-----+-----
Network | None
Exec   | Radius Start-Stop
System | Radius Stop-Only
```



Figure 43. Example of Optional Accounting Update Period and Accounting Suppression on Unknown User

Viewing RADIUS Statistics

General RADIUS

Syntax: show radius [host < ip-addr >] *Shows general RADIUS configuration, including the server IP addresses. Shows data for a specific RADIUS host. To use this command, the server's IP address must be configured in the switch.*

```

HP2512(config)# show radius
Status and Counters - General RADIUS Information
Deadtime(min) : 5
Timeout(secs) : 10
Retransmit Attempts : 2
Global Encryption Key : myg10balkey

      Auth  Acct
Server IP Addr  Port  Port  Encryption Key
-----
192.33.12.65   1812 1813  my65key

```

Figure 44. Example of General RADIUS Information from Show Radius Command

```

HP2512(config)# show radius host 192.33.12.65
Status and Counters - RADIUS Server Information
Server IP Addr : 192.33.12.65

Authentication UDP Port : 1812           Accounting UDP Port : 1813
Round Trip Time         : 2              Round Trip Time     : 7
Pending Requests        : 0              Pending Requests    : 0
Retransmissions         : 0              Retransmissions     : 0
Timeouts                : 0              Timeouts            : 0
Malformed Responses    : 0              Malformed Responses : 0
Bad Authenticators      : 0              Bad Authenticators  : 0
Unknown Types          : 0              Unknown Types       : 0
Packets Dropped         : 0              Packets Dropped     : 0
Access Requests        : 2              Accounting Requests : 2
Access Challenges       : 0              Accounting Responses : 2
Access Accepts         : 2
Access Rejects         : 0

```

Figure 45. Example of RADIUS Server Information From the Show Radius Host Command

Term	Definition
Round Trip Time	The time interval between the most recent Accounting-Response and the Accounting-Request that matched it from this RADIUS accounting server.
PendingRequests	The number of RADIUS Accounting-Request packets sent to this server that have not yet timed out or received a response. This variable is incremented when an accounting-Request is sent and decremented due to receipt of an Accounting-Response, a timeout or a retransmission.
Retransmissions	The number of RADIUS Accounting-Request packets retransmitted to this RADIUS accounting server. Retransmissions include retries where the Identifier and Acct-Delay have been updated, as well as those in which they remain the same.
Timeouts	The number of accounting timeouts to this server. After a timeout the client may retry to the same server, send to a different server, or give up. A retry to the same server is counted as a retransmit as well as a timeout. A send to a different server is counted as an Accounting-Request as well as a timeout.
Malformed Responses	The number of malformed RADIUS Accounting-Response packets received from this server. Malformed packets include packets with an invalid length. Bad authenticators and unknown types are not included as malformed accounting responses.
Bad Authenticators	The number of RADIUS Accounting-Response packets which contained invalid authenticators received from this server.
Unknown Types	The number of RADIUS packets of unknown type which were received from this server on the accounting port.
Packets Dropped	The number of RADIUS packets which were received from this server on the accounting port and dropped for some other reason.
Requests	The number of RADIUS Accounting-Request packets sent. This does not include retransmissions.
AccessChallenges	The number of RADIUS Access-Challenge packets (valid or invalid) received from this server.
AccessAccepts	The number of RADIUS Access-Accept packets (valid or invalid) received from this server.
AccessRejects	The number of RADIUS Access-Reject packets (valid or invalid) received from this server.
Responses	The number of RADIUS packets received on the accounting port from this server.

RADIUS Authentication

Syntax: show authentication

Displays the primary and secondary authentication methods configured for the Console, Telnet, Port-Access (802.1x), and SSH methods of accessing the switch. Also displays the number of access attempts currently allowed in a session.

show radius authentication

*Displays NAS identifier and data on the configured RADIUS server and the switch's interactions with this server. (Requires prior use of the **radius-server host** command to configure a RADIUS server IP address in the switch. See "Configuring RADIUS Accounting" on page 102.)*

```
HP2512(config)# show authentication
Status and Counters - Authentication Information
Login Attempts : 2

      | Login      Login      Enable   Enable
Access Task | Primary    Secondary Primary   Secondary
-----+-----
Console   | Local      None      Local    None
Telnet    | Radius     Local     Radius   Local
Port-Access | Local
SSH       | Radius     Local     Radius   Local
```

Figure 46. Example of Authentication Information from the Show Authentication Command

```
HP2512(config)# show radius authentication

Status and Counters - RADIUS Authentication Information

NAS Identifier : HP2512
Invalid Server Addresses : 0

      UDP
Server IP Addr | Port  Timeouts  Requests  Challenges  Accepts  Rejects
-----|-----
192.33.12.65  | 1812  0          2          0           2         0
```

Figure 47. Example of RADIUS Authentication Information from a Specific Server

RADIUS Accounting

- Syntax:** show accounting *Lists configured accounting interval, "Empty User" suppression status, accounting types, methods, and modes.*
- show radius accounting *Lists accounting statistics for the RADIUS server(s) configured in the switch (using the **radius-server host** command).*
- show accounting sessions *Lists the accounting sessions currently active on the switch.*

```

HP2512(config)# show accounting

Status and Counters - Accounting Information

Interval(min) : 5
Suppress Empty User : Yes

Type      | Method Mode
-----+-----
Network  | None
Exec     | Radius Start-Stop
System   | Radius Stop-Only
    
```

Figure 48. Example of the Accounting Configuration in the Switch

```

HP2512(config)# show radius accounting

Status and Counters - RADIUS Accounting Information

NAS Identifier : HP2512
Invalid Server Addresses : 0

          UDP
Server IP Addr  Port  Timeouts  Requests  Responses
-----
192.33.12.65   1813  0          1          1
    
```

Figure 49. Example of RADIUS Accounting Information for a Specific Server

```

HP2512(config)# show accounting sessions

Active Accounted actions on CONSOLE, User radius Priv 2,
Session ID 1, EXEC Accounting record, 00:02:32 Elapsed
    
```

Figure 50. Example Listing of Active RADIUS Accounting Sessions on the Switch

Changing RADIUS-Server Access Order

The switch tries to access RADIUS servers according to the order in which their IP addresses are listed by the **show radius** command. Also, *when you add a new server IP address, it is placed in the highest empty position in the list.*

Adding or deleting a RADIUS server IP address leaves an empty position, but does not change the position of any other server addresses in the list. For example if you initially configure three server addresses, they are listed in the order in which you entered them. However, if you subsequently remove the second server address in the list and add a new server address, the new address will be placed second in the list.

Thus, to move a server address up in the list, you must delete it from the list, ensure that the position to which you want to move it is vacant, and then re-enter it. For example, suppose you have already configured the following three RADIUS server IP addresses in the switch:

```
HP2512(config)# show radius
Status and Counters - General RADIUS Information
  Deadtime(min) : 0
  Timeout(secs) : 5
  Retransmit Attempts : 3
  Global Encryption Key :
  Server IP Addr  Auth  Acct  Encryption Key
  -----
  10.10.10.1     1812 1813
  10.10.10.2     1812 1813
  10.10.10.3     1812 1813
```

RADIUS server IP addresses listed in the order in which the switch will try to access them. In this case, the server at IP address 1.1.1.1 is first.

Note: If the switch successfully accesses the first server, it does not try to access any other servers in the list, even if the client is denied access by the first server.

Figure 51. Search Order for Accessing a RADIUS Server

To exchange the positions of the addresses so that the server at 10.10.10.003 will be the first choice and the server at 10.10.10.001 will be the last, you would do the following:

1. Delete 10.10.10.003 from the list. This opens the third (lowest) position in the list.
2. Delete 10.10.10.001 from the list. This opens the first (highest) position in the list.
3. Re-enter 10.10.10.003. Because the switch places a newly entered address in the highest-available position, this address becomes first in the list.
4. Re-enter 10.10.10.001. Because the only position open is the third position, this address becomes last in the list.

```

HP2512(config)# no radius host 10.10.10.003
HP2512(config)# no radius host 10.10.10.001
HP2512(config)# radius host 10.10.10.003
HP2512(config)# radius host 10.10.10.001

HP2512(config)# show radius

Status and Counters - General RADIUS Information
Deadtime(min) : 0
Timeout(secs) : 5
Retransmit Attempts : 3
Global Encryption Key :

Server IP Addr      Auth  Acct
                   Port  Port  Encryption Key
-----
10.10.10.3          1812 1813
10.10.10.2          1812 1813
10.10.10.1          1812 1813
    
```

Removes the "003" and "001" addresses from the RADIUS server list.

Inserts the "003" address in the first position in the RADIUS server list, and inserts the "001" address in the last position in the list.

Shows the new order in which the switch searches for a RADIUS server.

Figure 52. Example of New RADIUS Server Search Order

Messages Related to RADIUS Operation

Message	Meaning
Can't reach RADIUS server < x.x.x.x >.	A designated RADIUS server is not responding to an authentication request. Try pinging the server to determine whether it is accessible to the switch. If the server is accessible, then verify that the switch is using the correct encryption key and that the server is correctly configured to receive an authentication request from the switch.
No server(s) responding.	The switch is configured for and attempting RADIUS authentication, however it is not receiving a response from a RADIUS server. Ensure that the switch is configured to access at least one RADIUS server. (Use show radius .) If you also see the message Can't reach RADIUS server < x.x.x.x >, try the suggestions listed for that message.
Not legal combination of authentication methods.	Indicates an attempt to configure local as both the primary and secondary authentication methods. If local is the primary method, then none must be the secondary method.

Troubleshooting RADIUS Operation

Symptom	Possible Cause
<p>The switch does not receive a response to RADIUS authentication requests. In this case, the switch will attempt authentication using the secondary method configured for the type of access you are using (console, Telnet, or SSH).</p> <p>RADIUS server fails to respond to a request for service, even though the server's IP address is correctly configured in the switch.</p>	<p>There can be several reasons for not receiving a response to an authentication request. Do the following:</p> <ul style="list-style-type: none"> • 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. <p>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, then 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, then it overrides the global key and must match the server key.</p>

```

10.33.18.119(config)# show radius
Status and Counters - General RADIUS Information
  Deadtime(min) : 0
  Timeout(secs) : 5
  Retransmit Attempts : 3
  Global Encryption Key (My-Global-Key)

          Auth  Acct
Server IP Addr  Port  Port  Encryption Key
-----
10.33.18.119   1812  1813  119-only-key
  
```

Global RADIUS Encryption Key

Unique RADIUS Encryption Key for the RADIUS server at 10.33.18.119

— *This page is intentionally unused.* —

IP Preserve: Retaining VLAN-1 IP Addressing Across - Configuration File Downloads-

IP Preserve enables you to copy a configuration file to multiple Series 2500 switches while retaining the individual IP address and subnet mask on VLAN 1 in each switch, and the Gateway IP address assigned to the switch. This enables you to distribute the same configuration file to multiple switches without overwriting their individual IP addresses.

Operating Rules for IP Preserve

When **ip preserve** is entered as the last line in a configuration file stored on a TFTP server:

- If the switch's current IP address for VLAN 1 was not configured by DHCP/Bootp, IP Preserve retains the switch's current IP address, subnet mask, and IP gateway address when the switch downloads the file and reboots. The switch adopts all other configuration parameters in the configuration file into the startup-config file.
- If the switch's current IP addressing for VLAN 1 is from a DHCP server, IP Preserve is suspended. In this case, whatever IP addressing the configuration file specifies is implemented when the switch downloads the file and reboots. If the file includes DHCP/Bootp as the IP addressing source for VLAN 1, the switch will configure itself accordingly and use DHCP/Bootp. If instead, the file includes a dedicated IP address and subnet mask for VLAN 1 and a specific gateway IP address, then the switch will implement these settings in the startup-config file.
- The **ip preserve** statement does not appear in **show config** listings. To verify IP Preserve in a configuration file, open the file in a text editor and view the last line. For an example of implementing IP Preserve in a configuration file, see figure 53, below.

To set up IP Preserve, enter the **ip preserve** statement at the end of a configuration file. (Note that you do not execute IP Preserve by entering a command from the CLI).

```

; J4812A Configuration Editor; Created on release #F.05.17
hostname "HP2512"
time daylight-time-rule None
cdp run
.
.
.
password manager
password operator
ip preserve

```

Entering "ip preserve" in the last line of a configuration file implements IP Preserve when the file is downloaded to the switch and the switch reboots.

Figure 53. Example of Implementing IP Preserve in a Configuration File

For example, consider Figure 54:

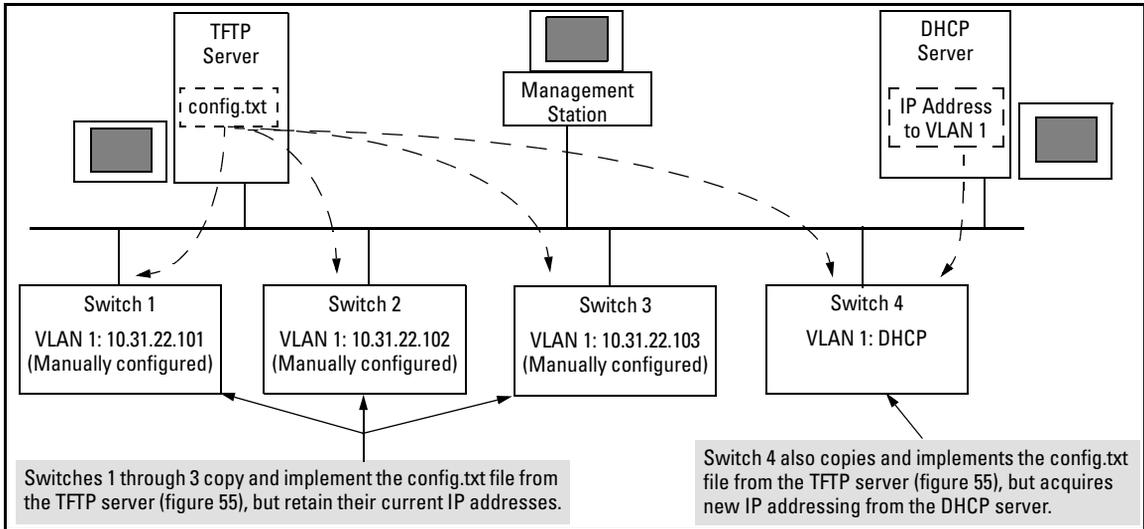


Figure 54. Example of IP Preserve Operation

If you apply the following configuration file to Figure 54, switches 1 - 3 will retain their manually assigned IP addressing and switch 4 will be configured to acquire its IP addressing from a DHCP server.

```

; J4812A Configuration Editor; Created on release #F.05.17
hostname "HP2512"
time daylight-time-rule None
cdp run
interface 11
    no lACP
exit
interface 12
    no lACP
exit
trunk 11-12 Trk1 Trunk
ip default-gateway 10.33.32.1
snmp-server community "public" Unrestricted
vlan 1
    name "DEFAULT_VLAN"
    ip address dhcp-bootp
    exit
password manager
password operator
ip preserve
    
```

Using figure 54, above, switches 1 - 3 ignore these entries because the file implements IP Preserve and their current IP addressing was not acquired through DHCP/Bootp.

Switch 4 ignores IP Preserve and implements the DHCP/Bootp addressing and IP Gateway specified in this file (because its last IP addressing was acquired from a DHCP/Bootp server).

IP Preserve Command

Figure 55. Configuration File in TFTP Server, with DHCP/Bootp Specified as the IP Addressing Source

If you apply this configuration file to figure 54, switches 1 - 3 will still retain their manually assigned IP addressing. However, switch 4 will be configured with the IP addressing included in the file.

```

; J4812A Configuration Editor; Created on release #F.05.17

hostname "HP2512"
time daylight-time-rule None
cdp run
interface 11
  no lACP
exit
interface 12
  no lACP
exit
trunk 11-12 Trk1 Trunk
ip default-gateway 10.33.32.1
snmp-server community "public" Unrestricted
vlan 1
  name "DEFAULT_VLAN"
  forbid 3
  untagged 1,7-10,13-14,Trk1
  tagged 4-6
  no untagged 2-3
  ip address 10.31.22.255 255.255.248.0
  exit
password manager
password operator
ip preserve

```

Because switch 4 (figure 54) received its most recent IP addressing from a DHCP/Bootp server, the switch ignores the **ip preserve** command and implements the IP addressing included in this file.

Figure 56. Configuration File in TFTP Server, with Dedicated IP Addressing Instead of DHCP/Bootp

To summarize the IP Preserve effect on IP addressing:

- If the switch received its most recent VLAN 1 IP addressing from a DHCP/Bootp server, it ignores the IP Preserve command when it downloads the configuration file, and implements whatever IP addressing instructions are in the configuration file.
- If the switch did not receive its most recent VLAN 1 IP addressing from a DHCP/Bootp server, it retains its current IP addressing when it downloads the configuration file.
- The content of the downloaded configuration file determines the IP addresses and subnet masks for other VLANs.

Configuring Port-Based Priority for Incoming Packets

Feature	Default	Menu	CLI	Web
Assigning a priority level to traffic on the basis of incoming port	Disabled	n/a	page 122	n/a

When network congestion occurs, it is important to move traffic on the basis of relative importance. However, without prioritization:

- Traffic from less important sources can consume bandwidth and slow down or halt delivery of more important traffic.
- Most traffic from all ports is forwarded as normal priority, and competes for bandwidth with all other normal-priority traffic, regardless of its relative importance.

Traffic received in tagged VLAN packets carries a specific 802.1p priority level (0 - 7) that the switch recognizes and uses to assign packet priority at the outbound port. With the default port-based priority, the switch handles traffic received in untagged packets as "Normal" (priority level = 0).

You can assign a priority level to inbound, untagged VLAN packets. (The switch does not alter the priority level of 802.1p tagged VLAN packets it receives.) Thus, for example, high-priority tagged VLAN traffic received on a port retains its priority in the switch. However, you have the option of configuring the port to assign a priority level to untagged VLAN traffic the port receives.

The Role of 802.1Q VLAN Tagging

An 802.1Q-tagged VLAN packet carries the packet's VLAN assignment and the 802.1p priority setting (0 - 7). (By contrast, an untagged packet does not have a tag and does not carry a priority setting.) Generally, the switch preserves and uses a packet's priority setting to determine which outbound queue the packet belongs in on the outbound port. If the outbound port is a tagged member of the VLAN, the packet carries its original priority to the next, downstream device. If the outbound port is not configured as a tagged member of the VLAN, then the tag is stripped from the packet, which then exits from the switch without a priority setting.

Outbound Port Queues and Packet Priority Settings

Series 2500 switch ports use two outbound port queues, *Normal* and *High*. As described below, these two queues map to the eight priority settings specified in the 802.1p standard.

Table 8. Mapping Priority Settings to Device Queues

802.1p Priority Settings Used In Tagged VLAN Packets	Series 2500 Outbound Port Queues	Queue Assignment in Downstream Devices With:			
		8 Queues	4 Queues	3 Queues*	2 Queues
1 (low)	Normal	1	1	1	1
2 (low)	Normal	2	1	1	1
0 (normal priority)	Normal	3	2	2	1
3	Normal	4	2	2	1
4	High	5	3	3	2
5	High	6	3	3	2
6	High	7	4	3	2
7 (high priority)	High	8	4	3	2

* HP ProCurve Switch 4100GL ports use three outbound priority queues.

For example, suppose you have configured port 10 to assign a priority level of 1 (low) to the (untagged) inbound packets it receives:

- An untagged packet coming into the switch on port 10 and leaving the switch through any other port configured as a tagged VLAN member would leave the switch as a tagged packet with a priority level of 1.
- A tagged packet with any 802.1p priority setting (0 - 7) coming into the switch on port 10 and leaving the switch through any other port configured as a tagged VLAN member would keep its original priority setting (regardless of the port-based priority setting on port 10).

Note

For a packet to carry a given 802.1p priority level from end-to-end in a network, the VLAN for the packet must be configured as tagged on all switch-to-switch links. Otherwise the tag is removed and the 802.1p priority is lost as the packet moves from one switch to the next.

Operating Rules for Port-Based Priority on Series 2500 Switches

- In the switch's default configuration, port-based priority is configured as "0" (zero) for inbound traffic on all ports.
- On a given port, when port-based priority is configured as "0" (zero) or 1 - 7, an inbound, *untagged* packet adopts the specified priority and is sent to the corresponding outbound queue on the outbound port. (See table 8, "Mapping Priority Settings to Device Queues", on page 121.) If the outbound port is a tagged member of the applicable VLAN, then the packet carries a tag with that priority setting to the next downstream device.
- On a given port, an inbound, *tagged* packet received on the port keeps the priority specified in the tag and is assigned an outbound queue on the basis of that priority (regardless of the port-based priority configured on the port). (Refer to table 8, "Mapping Priority Settings to Device Queues" on page 121.)
- If a packet leaves the switch through an outbound port configured as an untagged member of the packet's VLAN, then the packet leaves the switch without a VLAN tag and thus without an 802.1p priority setting.
- Trunked ports do not allow non-default (1 - 7) port-based priority settings. If you configure a non-default port-based priority value on a port and then add the port to a port trunk, then the port-based priority for that port is returned to the default "0".

Configuring and Viewing Port-Based Priority

This command enables or disables port-based priority on a per-port basis. You can either enter the command on the interface context level or include the interface in the command.

Syntax: qos priority < 1 - 7 >	<i>Configures a non-default port-based 802.1p priority for incoming, untagged packets on the designated ports, as described under "Operating Rules for Port-Based Priority", above.</i>
qos priority 0	<i>Returns a port-based priority setting to the default "0" for untagged packets received on the designated port(s). In this state the switch handles the untagged packets with "Normal" priority. (Refer to Table 8 on page 121.)</i>
show running-config	<i>Lists any non-default (1 - 7) port-based priority settings in the running-config file on a per-port basis.</i>
show config	<i>Lists any non-default (1 - 7) port-based priority settings in the startup-config file on a per-port basis. If the priority is set to the (default) "0", the setting is not included in the show config listing.</i>

For example, suppose you wanted to configure ports 10-12 on the switch to prioritize all untagged, inbound VLAN traffic as "Low" (priority level = 1; refer to table 8 on page 121).

```

HP2512(config)# interface e 9-12 qos priority 1
HP2512(config)# write mem
HP2512(config)# show config
Startup configuration:

: J4812A Configuration Editor; Created on release #F.05.17

hostname "HP2512"
time daylight-time-rule None
_cdp_run_ _ _ _ _
| interface 9 |
|   qos priority 1 |
| exit |
| interface 10 |
|   qos priority 1 |
| exit |
| interface 11 |
|   qos priority 1 |
| exit |
| interface 12 |
|   qos priority 1 |
| exit |
|_exit_|
snmp-server community "public" Unrestricted
vlan 1
  name "DEFAULT_VLAN"
-- MORE --, next page: Space, next line: Enter, quit: Control-C

```

Configures port-based priority on ports 9-12 to "1" (Low) and saves the configuration changes to the startup-config file.

Ports 9-12 are now configured to assign a priority level of "1" (Low) to untagged, incoming traffic. (Any inbound, tagged traffic retains its priority level while transiting the switch.)

Figure 57. Example of Configuring Non-Default Prioritization on Untagged, Inbound Traffic

Messages Related to Prioritization

Message	Meaning
< priority-level >: Unable to create.	The port(s) on which you are trying to configure a qos priority may belong to a port trunk. Trunked ports cannot be configured for qos priority.

Troubleshooting Prioritization

Symptom	Possible Cause
Ports configured for non-default prioritization (level 1 - 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.

Using the "Kill" Command To Terminate Remote Sessions

Using the **kill** command, you can terminate remote management sessions. (**Kill** does not terminate a Console session on the serial port, either through a direct connection or via a modem.)

Syntax: kill [<session-number>]

For example, if you are using the switch's serial port for a console session and want to terminate a currently active Telnet session, you would do the following:

```

HP2512(config)# show ip ssh
  SSH Enabled           : Yes

  IP Port Number       : 22
  Timeout (sec)       : 120
  Server Key Size (bits) : 512

  Ses Type      Source IP and Port
  -----
  1  console
  2  telnet
  3  ssh       15.30.252.195:1531
  4  inactive

HP2512(config)# kill 2
HP2512(config)# show ip ssh
  SSH Enabled           : Yes

  IP Port Number       : 22
  Timeout (sec)       : 120
  Server Key Size (bits) : 512

  Ses Type      Source IP and Port
  -----
  1  console
  2  inactive
  3  ssh       15.30.252.195:1531
  4  inactive
  
```

Figure 58. Example of Using the "Kill" Command To Terminate a Remote Session

Configuring Rapid Reconfiguration Spanning Tree (RSTP)

This section is related to the information on “Spanning Tree Protocol” in your *Series 2500 Switches Management and Configuration Guide* (5969-2354), but it primarily describes the new information associated with the new Spanning Tree standard, IEEE 802.1w (RSTP), which is supported by the F.04.08 release of your switch software.

You are referred to the *Management and Configuration Guide* for general information on the operation of Spanning Tree and for information on the older version of Spanning Tree, IEEE 802.1d (STP), which the F.04.08 software continues to support.

Overview

RSTP Feature	Default	Menu	CLI	Web
Viewing the RSTP/STP configuration	--	page 134	page 128	n/a
enable/disable RSTP/STP (RSTP is selected as the default protocol)	disabled	page 134	page 129	page 135
reconfiguring whole-switch values	Protocol Version: RSTP Force Version: RSTP-operation Switch Priority: step 8 Hello Time: 2 seconds Max Age: 20 seconds Forward Delay: 15 seconds	page 134	page 130	n/a
reconfiguring per-port values	Path Cost: depends on port type Priority: step 8 Edge Port: Yes Point-to-point: Force-true MCheck: Yes	page 134	page 132	n/a

As indicated in the manual, the Spanning Tree Protocol is used to ensure that only one active path at a time exists between any two end nodes in the network in which your switch is installed. Multiple paths cause a loop in the network over which broadcast and multicast messages are repeated continuously, which floods the network with traffic creating a broadcast storm.

In networks where there is more than one physical path between any two nodes, enabling Spanning Tree ensures a single active path between two such nodes by selecting the one most efficient path and blocking the other redundant paths. If a switch or bridge in the path becomes disabled, Spanning Tree activates the necessary blocked segments to create the next most efficient path.

The IEEE 802.1d version of Spanning Tree (STP) can take a fairly long time to resolve all the possible paths and to select the most efficient path through the network. The IEEE 802.1w Rapid Reconfiguration Spanning Tree (RSTP) significantly reduces the amount of time it takes to establish the network path. The result is reduced network downtime and improved network robustness.

In addition to faster network reconfiguration, RSTP also implements greater ranges for port path costs to accommodate the higher and higher connection speeds that are being implemented.

Transitioning from STP to RSTP

IEEE 802.1w RSTP is designed to be compatible with IEEE 802.1d STP. Even if all the other devices in your network are using STP, you can enable RSTP on your switch, and even using the default configuration values, your switch will interoperate effectively with the STP devices. If any of the switch ports are connected to switches or bridges on your network that do not support RSTP, RSTP can still be used on this switch. RSTP automatically detects when the switch ports are connected to non-RSTP devices in the Spanning Tree and communicates with those devices using 802.1d STP BPDU packets.

Because RSTP is so much more efficient at establishing the network path, though, that it is highly recommended that all your network devices be updated to support RSTP. RSTP offers convergence times of less than one second under optimal circumstances. To make the best use of RSTP and achieve the fastest possible convergence times, though, there are some changes that you should make to the RSTP default configuration. See “Optimizing the RSTP Configuration” below, for more information on these changes.

Note

Under some circumstances, it is possible for the rapid state transitions employed by RSTP to result in an increase in the rates of frame duplication and misordering in the switched LAN. In order to allow RSTP switches to support applications and protocols that may be sensitive to frame duplication and misordering, setting the Force Protocol Version parameter to STP-compatible allows RSTP to be operated with the rapid transitions disabled. The value of this parameter applies to all ports on the switch. See information on Force Version on page 130.

As indicated above, one of the benefits of RSTP is the implementation of a larger range of port path costs, which accommodates higher network speeds. New default values have also been implemented for the path costs associated with the different network speeds. This can create some incompatibility between devices running the older 802.1d STP and your switch running RSTP. Please see the “Note on Path Cost” on page 133 for more information on adjusting to this incompatibility.

Configuring RSTP

The default switch configuration has Spanning Tree disabled with RSTP as the selected protocol. That is, when Spanning Tree is enabled, RSTP is the version of Spanning Tree that is enabled, by default.

Optimizing the RSTP Configuration

To optimize the RSTP configuration on your switch, follow these steps (note that for the **Menu** method, all of these steps can be performed at the same time by making all the necessary edits on the Spanning Tree Operation screen and then saving the configuration changes):

1. Set the switch to support RSTP (RSTP is the default):

CLI: spanning-tree protocol-version rstp

Menu: Main Menu → 2. Switch Configuration → 4. Spanning Tree Operation → select Protocol Version: RSTP

2. Set the “point-to-point-mac” value to false on all ports that are connected to shared LAN segments (that is, to connections to hubs):

CLI: spanning-tree [ethernet] <port-list> point-to-point-mac force-false

Menu: Main Menu → 2. Switch Configuration → 4. Spanning Tree Operation → for each appropriate port, select Point-to-Point: Force-False

3. Set the “edge-port” value to false for all ports connected to other switches, bridges, and hubs:

CLI: no spanning-tree [ethernet] <port-list> edge-port

Menu: Main Menu → 2. Switch Configuration → 4. Spanning Tree Operation → for each appropriate port, select Edge: No

4. Set the “mcheck” value to false for all ports that are connected to devices that are known to be running IEEE 802.1d STP:

CLI: no spanning-tree [ethernet] <port-list> mcheck

Menu: Main Menu → 2. Switch Configuration → 4. Spanning Tree Operation → for each appropriate port, select MCheck: No

5. Enable RSTP Spanning Tree:

CLI: spanning-tree

Menu: Main Menu → 2. Switch Configuration → 4. Spanning Tree Operation → select STP Enabled: Yes

CLI: Configuring RSTP

Spanning Tree Commands in This Section	Applicable Protocol Version	Location
show spanning-tree config	both	Below on this page
spanning-tree	both	page 129
protocol-version <rstp stp>	both	page 129
force-version <rstp-operation stp-compatible>	RSTP	page 130
forward-delay <4 - 30>	both	page 130
hello-time <1 - 10>	both	page 130
maximum-age <6 - 40>	both	page 130
priority <0 - 15 0 - 65535>	RSTP STP	page 130
<[ethernet] port-list>	both	page 131
path-cost <1 - 200 000 000>	both	page 132
priority <0 - 15 0 - 65535>	RSTP STP	page 132
edge-port	RSTP	page 132
point-to-point-mac	RSTP	page 132
mcheck	RSTP	page 132
mode <norm fast>	STP	See the <i>Switch Management and Configuration Guide</i> for information on STP.
show spanning-tree	This command lists additional RSTP/STP monitoring data that is not covered in this section. See "Spanning Tree Protocol Information" in the "Monitoring and Analyzing Switch Operation" chapter in your <i>Switch Management and Configuration Guide</i> .	

Viewing the Current Spanning Tree Configuration. Even if Spanning Tree is disabled (the default configuration), the show spanning-tree config command lists the switch's full Spanning Tree configuration, including whole-switch and per-port settings.

Syntax: show spanning-tree configuration

Abbreviation: sho span config

In the default configuration, the output from this command appears similar to the following:

```

Spanning Tree Operation

Protocol Version : RSTP
STP Enabled [No] : Yes
Force Version [RSTP-operation] : RSTP-operation
Switch Priority [8] : 8           Hello Time [2] : 2
Max Age [20] : 20                Forward Delay [15] : 15

Port Type      | Cost      Priority Edge Point-to-Point MCheck
-----+-----
1  10/100TX | 200000    8      Yes Force-True   Yes
2  10/100TX | 200000    8      Yes Force-True   Yes
3  10/100TX | 200000    8      Yes Force-True   Yes
4  10/100TX | 200000    8      Yes Force-True   Yes
5  10/100TX | 200000    8      Yes Force-True   Yes
6  10/100TX | 200000    8      Yes Force-True   Yes
7  10/100TX | 200000    8      Yes Force-True   Yes
8  10/100TX | 200000    8      Yes Force-True   Yes
9  10/100TX | 200000    8      Yes Force-True   Yes
10 10/100TX | 200000    8      Yes Force-True   Yes
11 10/100TX | 200000    8      Yes Force-True   Yes
12 10/100TX | 200000    8      Yes Force-True   Yes
-- MORE --, next page: Space, next line: Enter, quit: Control-C

```

Figure 59. Example of the Spanning Tree Configuration Display

Enabling or Disabling RSTP. Issuing the command to enable Spanning Tree on the switch implements, by default, the RSTP version of Spanning Tree for all physical ports on the switch. Disabling Spanning Tree removes protection against redundant network paths.

Syntax: [no] spanning-tree

Abbreviation: [no] span

This command enables Spanning Tree with the current parameter settings or disables Spanning Tree, using the “no” option, without losing the most-recently configured parameter settings.

Enabling STP Instead of RSTP. If you decide, for whatever reason, that you would prefer to run the IEEE 802.1d (STP) version of Spanning Tree, then issue the following command:

Syntax: spanning-tree protocol-version stp

Abbreviation: span prot stp

For the STP version of Spanning Tree, the rest of the information in this section does not apply. Refer to the “Spanning Tree Protocol (STP)” section of your *Switch Management and Configuration Guide* for more information on the STP version and its parameters.

Reconfiguring Whole-Switch Spanning Tree Values. You can configure one or more of the following parameters, which affect the Spanning Tree operation of the whole switch:

Table 9. Whole-Switch RSTP Parameters

Parameter	Default	Description
protocol-version	RSTP	Identifies which of the Spanning Tree protocols will be used when Spanning Tree is enabled on the switch.
force-version	rstp-operation	<p>Sets the Spanning Tree compatibility mode. Even if rstp-operation is selected though, if the switch detects STP BPDU packets on a port, it will communicate to the attached device using STP BPDU packets.</p> <p>If errors are encountered, as described in the Note on page 126, the Force-Version value can be set to stp-compatible, which forces the switch to communicate out all ports using operations that are compatible with IEEE 802.1d STP.</p>
priority	32768 (8 as a step value)	<p>Specifies the protocol value used along with the switch MAC address to determine which device in the Spanning Tree is the root. The lower the priority value, the higher the priority.</p> <p>The value you enter has changed from the STP value. The range is 0 - 61440, but for RSTP the value is entered as a multiple (a step) of 4096. You enter a value in the range 0 - 15. The default value of 32768 is derived by the default setting of 8. Displaying the RSTP configuration (show spanning-tree config) shows 8, but displaying the RSTP operation (show spanning-tree) shows 32768.</p>
*maximum-age	20 seconds	Sets the maximum age of received Spanning Tree information before it is discarded. The range is 6 to 40 seconds.
*hello-time	2 seconds	Sets the time between transmission of Spanning Tree messages. Used only when this switch is the root. The range is 1 to 10 seconds.
*forward-delay	15 seconds	Sets the time the switch waits between transitioning ports from listening to learning and from learning to forwarding states. The range is 4 to 30 seconds.
<p>*These parameters are the same for RSTP as they are for STP. The switch uses its own maximum-age, hello-time, and forward-delay settings only if it is operating as the root device in the Spanning Tree. If another device is the root device, then the switch uses the other device's settings for these parameters.</p>		

Note

Executing the `spanning-tree` command alone enables Spanning Tree. Executing the command with one or more of the whole-switch RSTP parameters shown in the table on the previous page, or with any of the per-port RSTP parameters shown in the table on page 132, does not enable Spanning Tree. It only configures the Spanning Tree parameters, regardless of whether Spanning Tree is actually running (enabled) on the switch.

Using this facility, you can completely configure Spanning Tree the way you want and then enable it. This method minimizes the impact on the network operation.

Syntax:

```
spanning-tree
  protocol-version <rstp | stp>
  force-version <rstp-operation | stp-compatible>
  priority <0 - 15>
  maximum-age <6 - 40 seconds>
  hello-time <1- 10 seconds>
  forward-delay <4 - 30 seconds>
```

Abbreviations:

```
span
  prot <rstp | stp>
  forc <rstp | stp>
  pri <0 - 15>
  max <6 - 40>
  hello <1 - 10>
  forw <4 - 30>
```

Defaults: see the table on the previous page.

Multiple parameters can be included on the same command line. For example, to configure a maximum-age of 30 seconds and a hello-time of 3 seconds, you would issue the following command:

```
HP 2524 (config)# span max 30 hello 3
```

Reconfiguring Per-Port Spanning Tree Values. You can configure one or more of the following parameters, which affect the Spanning Tree operation of the specified ports only:

Table 10. Per-Port RSTP Parameters

Parameter	Default	Description
edge-port	Yes	Identifies ports that are connected to end nodes. During Spanning Tree establishment, these ports transition immediately to the Forwarding state. In this way, the ports operate very similarly to ports that are configured in “fast mode” under the STP implementation in previous HP switch software. Disable this feature on all switch ports that are connected to another switch, or bridge, or hub. Use the “no” option on the spanning tree command to disable edge-port.
mcheck	Yes	Ports with mcheck set to true are forced to send out RSTP BPDUs for 3 seconds. This allows for switches that are running RSTP to establish their connection quickly and for switches running 802.1d STP to be identified. If the whole-switch parameter Force-Version is set to “stp-compatible”, the mcheck setting is ignored and STP BPDUs are sent out all ports. Disable this feature on all ports that are known to be connected to devices that are running 802.1d STP. Use the “no” option on the spanning tree command to disable mcheck.
path-cost	10 Mbps – 2 000 000 100 Mbps – 200 000 1 Gbps – 20 000	Assigns an individual port cost that the switch uses to determine which ports are the forwarding ports. The range is 1 to 200,000,000 or auto. By default, this parameter is automatically determined by the port type, as shown by the different default values. If you have previously configured a specific value for this parameter, you can issue the command with the auto option to restore the automatic setting feature. Please see the Note on Path Cost on page 133 for information on compatibility with devices running 802.1d STP for the path cost values.
point-to-point-mac	force-true	This parameter is used to tell the port if it is connected to a point-to-point link, such as to another switch or bridge or to an end node (force-true). This parameter should be set to force-false for all ports that are connected to a hub, which is a shared LAN segment. You can also set this parameter to auto and the switch will automatically set the force-false value on all ports that it detects are not running at full duplex. All connections to hubs are not full duplex.
priority	128 (8 as a step value)	This parameter is used by RSTP to determine the port(s) to use for forwarding. The port with the lowest number has the highest priority. The range is 0 to 240, but you configure the value by entering a multiple of 16. You enter a value in the range 0 - 15. The default value of 128 is derived by the default setting of 8. Displaying the RSTP configuration (show spanning-tree config) shows 8, but displaying the RSTP operation (show spanning-tree) shows 128.

Syntax:

```
spanning-tree [ethernet] <port-list>
  path-cost <1 - 200000000>
  point-to-point-mac <force-true | force-false | auto>
  priority <0 - 15>

[no] spanning-tree [ethernet] <port-list>
  edge-port
  mcheck
```

Abbreviations:

```
span <port-list>
  path <1 - 200000000>
  forc <force-t | force-f | auto>
  pri <0 - 15>

[no] span <port-list>
  edge
  mch
```

Defaults: see the table on the previous page.

Note on Path Cost

RSTP implements a greater range of path costs and new default path cost values to account for higher network speeds. These values are different than the values defined by 802.1d STP as shown in the next table.

Port Type	802.1d STP Path Cost	RSTP Path Cost
10 Mbps	100	2 000 000
100 Mbps	10	200 000
1 Gbps	5	20 000
10 Gbps	?	2000

Because the maximum value for the path cost allowed by 802.1d STP is 65535, devices running that version of Spanning Tree cannot be configured to match the values defined by RSTP, at least for 10 Mbps and 100 Mbps ports. In LANs where there is a mix of devices running 802.1d STP and RSTP, you should reconfigure the devices so the path costs match for ports with the same network speeds.

Menu: Configuring RSTP

1. From the console CLI prompt, enter the menu command.
 HP ProCurve Switch # **menu**
2. From the switch console Main Menu, select
2. Switch Configuration ...
4. Spanning Tree Operation
3. Press [E] (for **Edit**) to highlight the **Protocol Version** parameter field.
4. Press the Space bar to select the version of Spanning Tree you wish to run: **RSTP** or **STP**.
Note: If you change the protocol version, you will have to reboot the switch for the change to take effect. See step 9 and step 10.
5. Press the [Tab] or down arrow key to go to the **STP Enabled** field. Note that when you do this, the remaining fields on the screen will then be appropriate for the version of Spanning Tree that was selected in step 3. The screen image below is for RSTP.
6. Press the Space bar to select **Yes** to enable Spanning Tree.

```

HP ProCurve Switch
===== TELNET - MANAGER MODE =====
                Switch Configuration - Spanning Tree Operation

Protocol Version : RSTP
STP Enabled [No] : No
Force Version [RSTP-operation] : RSTP-operation
Switch Priority [8] : 8                Hello Time [2] : 2
Max Age [20] : 20                    Forward Delay [15] : 15

Port   Type      Cost      Priority  Edge  Point-to-Point  MCheck
----   -+-----+-----+-----+-----+-----+-----+
 1    10/100TX | 200000    8        Yes   Force-True      Yes
 2    10/100TX | 200000    8        Yes   Force-True      Yes
 3    10/100TX | 200000    8        Yes   Force-True      Yes
 4    10/100TX | 200000    8        Yes   Force-True      Yes
 5    10/100TX | 200000    8        Yes   Force-True      Yes
 6    10/100TX | 200000    8        Yes   Force-True      Yes

Actions->  Cancel   Edit   Save   Help

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

Figure 60. Example of the RSTP Configuration Screen

7. Press the **[Tab]** key or use the arrow keys to go to the next parameter you want to change, then type in the new value or press the Space bar to select a value. (To get help on this screen, press **[Enter]** to select the **Actions** → line, then press **[H]**, for **Help**, to display the online help.)
8. Repeat step 6 for each additional parameter you want to change.
Please see “Optimizing the RSTP Configuration” on page 127 for recommendations on configuring RSTP to make it operate the most efficiently.
9. When you are finished editing parameters, press **[Enter]** to return to the **Actions** → line and press **[S]** to save the currently displayed Spanning Tree settings and return to the Main Menu.
10. If you have changed the Protocol Version, in step 1, reboot the switch now by selecting

6. Reboot Switch

Web: Enabling or Disabling RSTP

In the web browser interface, you can enable or disable Spanning Tree on the switch. If the default configuration is in effect such that RSTP is the selected protocol version, enabling Spanning Tree through the web browser interface will enable RSTP with its current configuration. To configure the other Spanning Tree features, telnet to the switch console and use the CLI or menu.

To enable or disable Spanning Tree using the web browser interface:

1. Click on the **Configuration** tab.
2. Click on **[Device Features]**.
3. Enable or disable Spanning Tree.
4. Click on **[Apply Changes]** to implement the configuration change.

— *This page is intentionally unused.* —

Enhancements in Release F.02.11

Enhancement	Summary	Page
Adds the fast-uplink spanning tree (STP) mode to spanning-tree operation	In an 802.1D STP environment with redundant links, an active link failure typically results in a convergence time of 30 seconds for a backup link to become the active, forwarding link. Fast-uplink STP reduces this time to approximately ten seconds.	below
Adds the show tech command to the switch troubleshooting capabilities	This command outputs, in a single listing, switch operating and running configuration details from several internal switch sources.	151

Fast-Uplink Spanning Tree Protocol (STP)

Fast-Uplink STP improves the recovery (convergence) time in wiring closet switches with redundant uplinks. Specifically, a Series 2500 switch having redundant links toward the root device can decrease the convergence time (or failover) to a new uplink (STP root) port to as little as ten seconds. To realize this performance, a Series 2500 switch must be:

- Used as a wiring closet switch (also termed an *edge switch* or a *leaf switch*).
- Configured for fast-uplink STP mode on two or more ports intended for redundancy in the direction of the root switch, so that at any time only one of the redundant ports is expected to be in the forwarding state.

Note

When properly implemented, fast-uplink STP offers a method for achieving faster failover times than standard STP. While fast-uplink STP remains an effective means for reducing failover time, HP recommends that you move to the Rapid Convergence STP (RSTP; 802.1w) available with software release F.04.08 and greater, for the best failover performance.

Caution

In general, fast-uplink spanning tree on the Series 2500 switches is useful when running STP in a tiered topology that has well-defined edge switches. Also, ensure that an interior switch is used for the root switch and for any logical backup root switches. You can accomplish this by using the Spanning Tree Priority (sometimes termed bridge priority) settings that define the primary STP root switch and at least one failover root switch (in the event that the primary root switch fails). Inappropriate use of Fast-Uplink STP can cause intermittent loops in a network topology. For this reason, the Fast-Uplink STP feature should be used only by experienced network administrators who have a strong understanding of the IEEE 802.1d standard and STP interactions and operation. If you want to learn more about STP operation, you may find it helpful to refer to publications such as:

Perlman, Radia, *Interconnections, Second Edition; Bridges, Routers, Switches, and Internet-working Protocols*, Addison-Wesley Professional Computing Series, October 1999

To use fast-uplink STP on a Series 2500 switch, configure fast-uplink (**Mode = Uplink**) only on the switch's upstream ports; (that is, two or more ports forming a group of redundant links in the direction of the STP root switch). If the active link in this group goes down, fast-uplink STP selects a different upstream port as the root port and resumes moving traffic in as little as ten seconds. The device(s) on the other end of the links must be running STP. However, because fast uplink should be configured only on the Series 2500 switch uplink ports, the device(s) on the other end of the links can be either HP devices or another vendor's devices, regardless of whether they support fast uplink. For example:

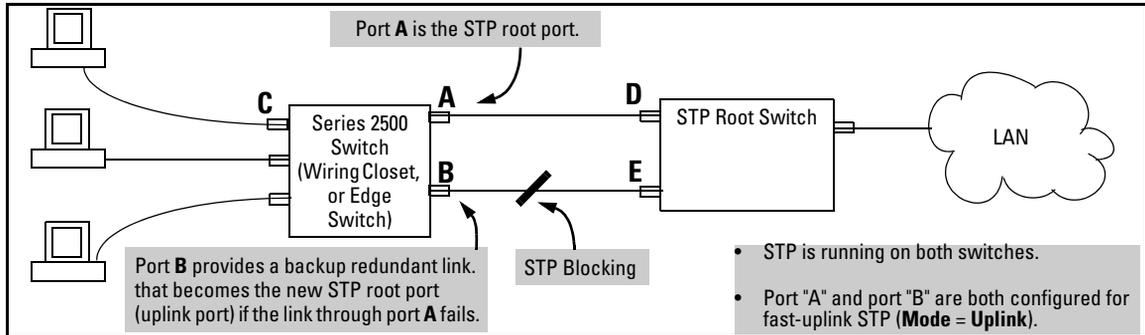


Figure 61. Example of How To Implement Fast-Uplink STP

Terminology

Term	Definition
downlink port (downstream port)	A switch port that is linked to a port on another switch (or to an end node) that is sequentially further away from the STP root device. For example, port "C" in figure 61, above, is a downlink port.
edge switch	For the purposes of fast-uplink STP, this is a switch that has no other switches connected to its downlink ports. An edge switch is sequentially further from the root device than other switches to which it is connected. Also termed <i>wiring closet switch</i> or <i>leaf switch</i> . For example, switch "4" in figure 62 (page 139) is an edge switch.
interior switch	In an STP environment, a switch that is sequentially closer to the STP root device than one or more other switches to which it is connected. For example, switches "1", "2", and "3" in figure 62 (page 139) are interior switches.
single-instance spanning tree	A single spanning-tree ensuring that there are no logical network loops associated with any of the connections to the switch, regardless of whether there are any VLANs configured on the switch. For more information, see "Spanning Tree Protocol (STP)" in chapter 9, "Configuring Advanced Features", in the Management and Configuration Guide for your Series 2500 switch.
uplink port (upstream port)	A switch port linked to a port on another switch that is sequentially closer to the STP root device. For example, ports "A" and "B" in figure 61 on page 138 are uplink ports.
wiring closet switch	Another term for an "edge" or "leaf" switch.

When single-instance spanning tree (STP) is running in a network and a forwarding port goes down, a blocked port typically requires a period of

$$(2 \times (\textit{forward delay}) + \textit{link down detection})$$

to transition to forwarding. In a normal spanning tree environment, this transition is usually 30 seconds (with the **Forward Delay** parameter set to its default of 15 seconds). However, by using the fast-uplink spanning tree feature, a port on a Switch 2512 or 2524 used as an *edge switch* can make this transition in as little as ten seconds. (In an STP environment, an *edge switch* is a switch that is connected only to switches that are closer to the STP root switch than the edge switch itself, as shown by switch "4" in figure 62, below.)

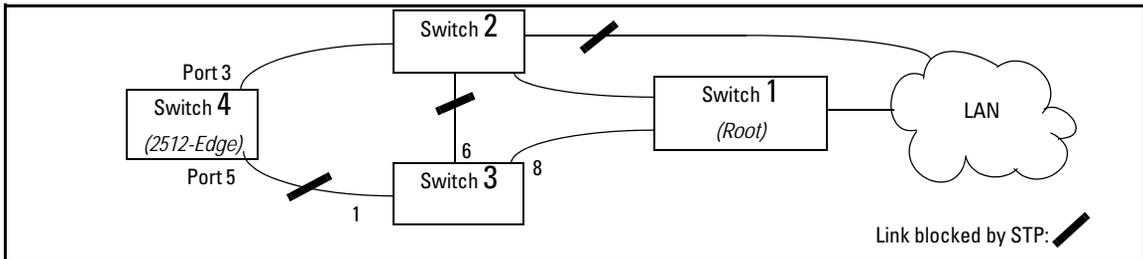


Figure 62. Example of an Edge Switch in a Topology Configured for STP Fast Uplink

In figure 62, STP is enabled and in its default configuration on all switches, unless otherwise indicated in table 11, below:

Table 11. STP Parameter Settings for Figure 62

STP Parameter	Switch "1"	Switch "2"	Switch "3"	Switch "4"
Switch Priority	0 ¹	1 ²	32,768 (default)	32,768 (default)
(Fast) Uplink	No	No	No	Ports 3 & 5

¹This setting ensures that Switch "1" will be the primary root switch for STP in figure 62.

²This setting ensures that Switch "2" will be the backup root switch for STP in figure 62.

With the above-indicated topology and configuration:

- **Scenario 1:** If the link between switches "4" and "2" goes down, then the link between switches "4" and "3" will begin forwarding in as little as ten seconds.
- **Scenario 2:** If Switch "1" fails, then:
 - Switch "2" becomes the root switch.
 - The link between Switch "3" and Switch "2" begins forwarding.
 - The link between Switch "2" and the LAN begins forwarding.

Operating Rules for Fast Uplink

- A switch with ports configured for fast uplink must be an edge switch and not either an interior switch or the STP root switch.

Configure fast-uplink on only the edge switch ports used for providing redundant STP uplink connections in a network. (Configuring Fast-Uplink STP on ports in interior switches can create network performance problems.) That is, a port configured for STP uplink should not be connected to a switch that is sequentially further away from the STP root device. For example, switch "4" in figure 62 (page 139) is an edge switch.

- Configure fast uplink on a group (two or more) of redundant edge-switch uplink ports where only one port in the group is expected to be in the forwarding state at any given time.
- Edge switches cannot be directly linked together using fast-uplink ports. For example, the connection between switches 4 and 5 in figure 63 is not allowed for fast-uplink operation.

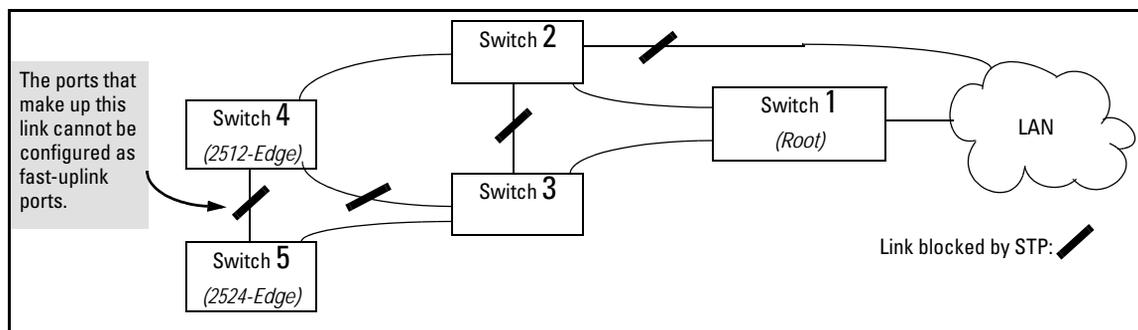


Figure 63. Example of a Disallowed Connection Between Edge Switches

- Apply fast-uplink only on the uplink ports of an edge switch. For example, on switch "4" (an edge switch) in figure 63 above, only the ports connecting switch "4" to switches "2" and "3" are upstream ports that would use fast uplink. Note also that fast uplink should *not* be configured on both ends of a point-to-point link, but only on the uplink port of an edge switch.
- Ensure that the switch you intend as a backup root device will in fact become the root if the primary root fails, and that no ports on the backup root device are configured for fast-uplink operation. For example, if the **STP Priority** is the same on all switches—default: 32768—then the switch with the lowest MAC address will become the root switch. If that switch fails, then the switch with the next-lowest MAC address will become the root switch. Thus, you can use **STP Priority** to control which switch STP selects as the root switch and which switch will become the root if the first switch fails.
- Fast-Uplink STP requires a minimum of two uplink ports.

Menu: Viewing and Configuring Fast-Uplink STP

You can use the menu to quickly display the entire STP configuration and to make any STP configuration changes.

To View and/or Configure Fast-Uplink STP. This procedure uses the Spanning Tree Operation screen to enable STP and to set the Mode for fast-uplink STP operation.

- From the Main Menu select:
 - 2. Switch Configuration ...**
 - 4. Spanning Tree Operation**
- In the default STP configuration, RSTP is the selected protocol version. If this is the case on your switch, you must change the Protocol Version to STP in order to use Fast-Uplink STP:

- If the **Protocol Version** is set to RSTP (the default, as shown in this example, go to step 3.
- If the **Protocol Version** is set to STP, the rest of the screen will appear as shown in figure 66. In this case, go to step 4 on page 143.

```

----- CONSOLE - MANAGER MODE -----
Switch Configuration - Spanning Tree Operation

Protocol Version : (RSTP)
STP Enabled [No] : No
Force Version [RSTP-operation] : RSTP-operation
Switch Priority [8] : 8           Hello Time [2] : 2
Max Age [20] : 20                Forward Delay [15] : 15

Port    Type          Cost    Priority  Edge  Point-to-Point  MCheck
----  -
3      10/100TX  | 200000    8        Yes   Force-True      Yes
4      10/100TX  | 200000    8        Yes   Force-True      Yes
5      10/100TX  | 200000    8        Yes   Force-True      Yes
6      10/100TX  | 200000    8        Yes   Force-True      Yes
7      10/100TX  | 200000    8        Yes   Force-True      Yes
8      10/100TX  | 200000    8        Yes   Force-True      Yes

Actions->  Cancel   Edit   Save   Help

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

Figure 64. The Default STP Screen With the Protocol Version Field Set to "RSTP"

3. If the Protocol Version is set to RSTP (as shown in figure 64), do the following:
 - a. Press **[E]** (**Edit**) to move the cursor to the **Protocol Version** field.
 - b. Press the Space bar once to change the **Protocol Version** field to STP.
 - c. Press **[Enter]** to return to the command line.
 - d. Press **[S]** (for **Save**) to save the change and exit from the Spanning Tree Operation screen. you will then see a screen with the following:

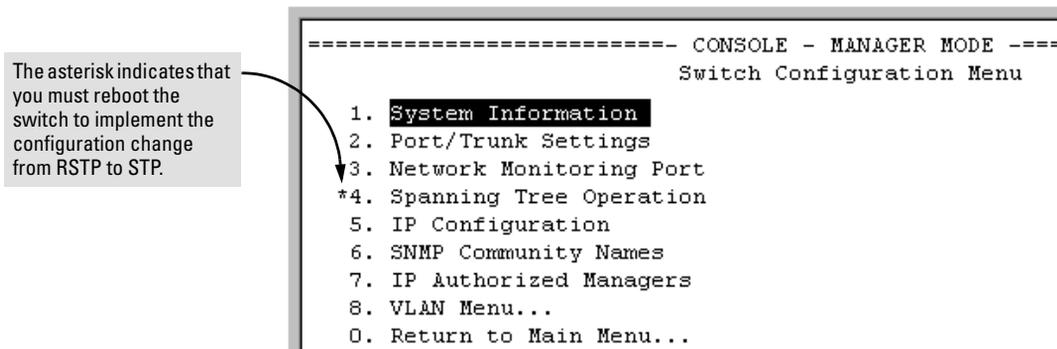


Figure 65. Changing from RSTP to STP Requires a System Reboot

- e. Press **[0]** (zero) to return to the Main Menu, then **[6]** to reboot the switch.
- f. After you reboot the switch, enter the menu command at the CLI to return to the Main Menu, then select:

- 2. Switch Configuration ...
- 4. Spanning Tree Operation

You will then see the Spanning-Tree screen with **STP** (802.1d) selected in the **Protocol Version** field (figure 66).

```

===== CONSOLE - MANAGER MODE =====
                Switch Configuration - Spanning Tree Operation
Protocol Version : STP
STP Enabled [No] : No
Switch Priority [32768] : 32768           Hello Time [2] : 2
Max Age [20] : 20                       Forward Delay [15] : 15

Port      Type          Cost      Priority  Mode
-----
A1 10/100TX | 100      128      Norm
A4 10/100TX | 100      128      Norm
A5 10/100TX | 100      128      Norm
A6 10/100TX | 100      128      Norm
A7 10/100TX | 100      128      Norm
A8 10/100TX | 100      128      Norm
A9 10/100TX | 100      128      Norm

Actions->  Cancel   Edit   Save   Help

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

```

In this example, ports 2 and 3 have already been configured as a port trunk (**Trk1**), which appears at the end of the port listing.

All ports (and the trunk) are in their default STP configuration.

Note: Ports 10-14 do not appear in this simulation. In the actual menu screen, you must scroll the cursor down the port list to view the trunk configuration.

Figure 66. The Spanning Tree Operation Screen

4. On the ports and/or trunks you want to use for redundant fast uplink connections, change the mode to **Uplink**. In this example, port 1 and Trk1 (using ports 2 and 3) provide the redundant uplinks for STP:
 - a. Press **[E]** (for **Edit**), then enable STP on the switch by using the Space bar to select **Yes** in the Spanning Tree Enabled field.
 - b. Use **[Tab]** to move to the Mode field for port 1.
 - c. Use the Space bar to select **Uplink** as the mode for port 1.
 - d. Use **[↓]** to move to the Mode field for Trk1.
 - e. Use the Space bar to select **Uplink** as the Mode for Trk1.
 - f. Press **[Enter]** to return the cursor to the Actions line.

```

===== CONSOLE - MANAGER MODE =====
                Switch Configuration - Spanning Tree Operation
Protocol Version : STP
STP Enabled [No] : Yes
Switch Priority [32768] : 32768      Hello Time [2] : 2
Max Age [20] : 20                  Forward Delay [15] : 15

Port      Type      Cost      Priority  Mode
-----
1         10/100TX  | 100      128      Uplink
4         10/100TX  | 100      128      Norm
5         10/100TX  | 100      128      Norm
.         .         | .         .         .
.         .         | .         .         .
.         .         | .         .         .
12        10/100TX  | 100      128      Norm
13        .         | 100      128      Norm
14        .         | 100      128      Norm
Trk1     .         | 100      64       Uplink

Actions->  _Cancel  Edit  _Save  _Help

```

STP is enabled.

Port 1 and Trk1 are now configured for fast-uplink STP.

Edit the fields displayed above.
Use arrow keys to change action selection and <Enter> to execute action.

Figure 67. Example of STP Enabled with Two Redundant Links Configured for Fast-Uplink STP

5. Press [S] (for **S**ave) to save the configuration changes to flash (non-volatile) memory.

To View Fast-Uplink STP Status. Continuing from figures 66 and 67 in the preceding procedure, this task uses the same screen that you would use to view STP status for other operating modes.

1. From the Main Menu, select:
 1. Status and Counters . . .
 7. Spanning Tree Information

```

===== CONSOLE - MANAGER MODE =====
                Status and Counters - Spanning Tree Information

STP Enabled           : Yes
Switch Priority       : 32,768
Hello Time           : 2
Max Age              : 20
Forward Delay        : 15

Topology Change Count : 2
Time Since Last Change : 15 mins

Root MAC Address     : 0060b0-889e00
Root Path Cost       : 20
Root Port            : Trk1
Root Priority         : 16000

Actions->  Back  Show ports  Help

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

Indicates which uplink is the active path to the STP root device.
Note: A switch using fast-uplink STP must never be the STP root device.

Figure 68. Example of STP Status with Trk1 (Trunk 1) as the Path to the STP Root Device

2. Press [S] (for **S**how ports) to display the status of individual ports.

```

===== CONSOLE - MANAGER MODE =====
                Status and Counters - Spanning Tree - Port Information

Port   Type      Cost  Priority  State      Designated Bridge
-----
1      10/100TX    10    128      Blocking   0030c1-7fcc40
4      10/100TX    10    128      Disabled
5      10/100TX    10    128      Forwarding 0030c1-a914c0
6      10/100TX    10    128      Forwarding 0030c1-a919c1
.      .           .     .         .
.      .           .     .         .
12     10/100TX    10    128      Forwarding 0030c1-c884c0
13     .           100   128      Disabled
14     .           100   128      Disabled
Trk1   .           10    64       Forwarding 0030c1-7fcc40

Actions->  Back  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.
    
```

Redundant STP Link in (Fast) Uplink Mode
 Links to PC or Workstation End Nodes
 Redundant STP Link in (Fast) Uplink Mode

Figure 69. Example of STP Port Status with Two Redundant STP Links

In figure 69:

- Port 1 and Trk1 (trunk 1; formed from ports 2 and 3) are redundant fast-uplink STP links, with trunk 1 forwarding (the active link) and port 1 blocking (the backup link). (To view the configuration for port 1 and Trk1, see figure 67 on page 144.)
- If the link provided by trunk 1 fails (on both ports), then port 1 begins forwarding in fast-uplink STP mode.
- Ports 5, 6, and 12 are connected to end nodes and do not form redundant links.

CLI: Viewing and Configuring Fast-Uplink STP

Using the CLI to View Fast-Uplink STP. You can view fast-uplink STP using the same **show** commands that you would use for standard STP operation:

Syntax: show spanning-tree *Lists STP status.*
 show spanning-tree config *Lists STP configuration for the switch and for individual ports.*

For example, figures 70 and 71 illustrate a possible topology, STP status listing, and STP configuration for a Series 2500 switch with:

- STP enabled and the switch operating as an Edge switch
- Port 1 and trunk 1 (Trk1) configured for fast-uplink STP operation
- Several other ports connected to PC or workstation end nodes

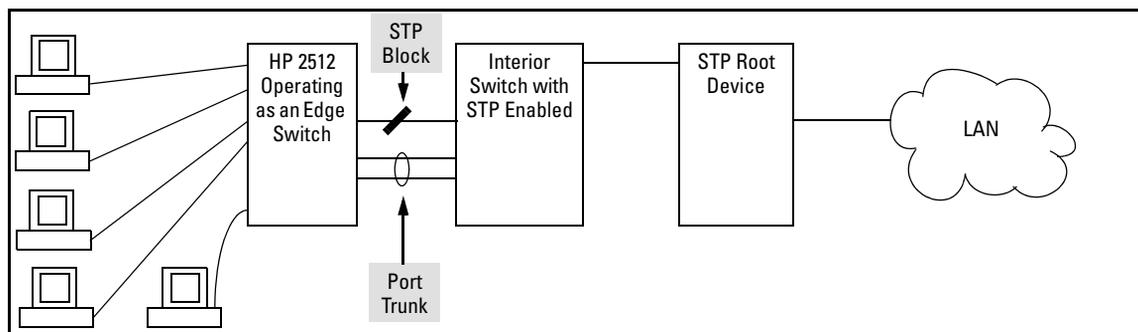


Figure 70. Example Topology for the Listing Shown in Figure 71

```

HP2512(config)# show spanning-tree_
Status and Counters - Spanning Tree Information

STP Enabled           : Yes
Switch Priority       : 32,768
Hello Time           : 2
Max Age              : 20
Forward Delay        : 15

Topology Change Count : 25
Time Since Last Change : 13 mins

Root MAC Address     : 0001e7-a09900
Root Path Cost       : 20
Root Port            : Trk1
Root Priority         : 16768

Port   Type      Cost  Priority  State      | Designated Bridge
-----+-----
 1    10/100TX   10    128     Blocking   | 0030c1-a9c800
 4    10/100TX   10    128     Disabled   |
 5    10/100TX   10    128     Forwarding | 0030c1-7fec40
 6    10/100TX   10    128     Forwarding | 0030c1-a9c800
-- MORE --
 7    10/100TX   10    128     Forwarding | 0030c1-a9c822
 8    10/100TX   10    128     Disabled   |
 9    10/100TX   10    128     Forwarding | 00a0c9-a234c3
10    10/100TX   10    128     Forwarding | 0030c1-449bc0
11    10/100TX   10    128     Disabled   |
12    10/100TX   10    128     Disabled   |
13           100    128     Disabled   |
14           100    128     Disabled   |
Trk1           10    64     Forwarding | 0030c1-a9c800
    
```

Indicates that Trk1 (Trunk 1) provides the currently active path to the STP root device.

Redundant STP link in the Blocking state.

Links to PC or Workstation End Nodes

Redundant STP link in the Forwarding state. (See the "Root Port" field, above. This is the currently active path to the STP root device.)

Figure 71. Example of a Show Spanning-Tree Listing for the Topology Shown in Figure 70

```

HP2512(config)# show spanning-tree config
Spanning Tree Operation
Spanning Tree Enabled : Yes
STP Priority : 32768
Max Age : 20
Hello Time : 2
Forward Delay : 15

Port Type      | Cost  Pri Mode
-----+-----+-----
1   10/100TX   | 10   128 Uplink
4   10/100TX   | 10   128 Norm
5   10/100TX   | 10   128 Norm
6   10/100TX   | 10   128 Norm
7   10/100TX   | 10   128 Norm
8   10/100TX   | 10   128 Norm
9   10/100TX   | 10   128 Norm
10  10/100TX   | 10   128 Norm
11  10/100TX   | 10   128 Norm
12  10/100TX   | 10   128 Norm
13          | 100  128 Norm
14          | 100  128 Norm
Trk1 Trunk    | 10   64  Uplink
    
```

Figure 72. Example of a Configuration Supporting the STP Topology Shown in Figure 70

Using the CLI To Configure Fast-Uplink STP. This example uses the CLI to configure the switch for the fast-uplink operation shown in figures 70, 71, and 72. (The example assumes that ports 2 and 3 are already configured as members of the port trunk—Trk1, and all other STP parameters are left in their default state.)

Note that the default STP Protocol Version is RSTP (Rapid STP, or 802.1w). Thus, if the switch is set to the STP default, you must change it to the STP (802.1d) Protocol Version before you can configure Fast-Uplink. For example:

```

HP2512(config)# show spanning-tree
Status and Counters - Spanning Tree Information
Protocol Version : RSTP
STP Enabled : No

Port Type      Cost      Priority State | Designated Bridge
-----+-----+-----+-----+-----
1   10/100TX   10       128    STP |
4   10/100TX   10       128    STP |
5   10/100TX   10       128    STP |
6   10/100TX   10       128    STP |
7   10/100TX   10       128    STP |
8   10/100TX   10       128    STP |
9   10/100TX   10       128    STP |
10  10/100TX   10       128    STP |
11  10/100TX   10       128    STP |
12  10/100TX   10       128    STP |
13          100      128    STP |
14          100      128    STP |
Trk1 Trunk    10       64     STP |

HP2512(config)# (spanning-tree protocol-version stp)
STP version was changed. To activate the change you must
save the configuration to flash and reboot the device.
HP2512(config)# (write mem)
HP2512(config)# (boot)
Device will be rebooted, do you want to continue [y/n]? y
    
```

Figure 73. Example of Changing the STP Configuration from the Default RSTP (802.1w) to STP (802.1d)

Syntax: spanning-tree e <port/trunk-list> mode uplink *Enables STP on the switch and configures fast-uplink STP on the designated interfaces (port or trunk).*

```
HP2512(config)# spanning-tree e 1,trk1 mode uplink
```

Operating Notes

Effect of Reboots on Fast-Uplink STP Operation. When configured, fast-uplink STP operates on the designated ports in a running Series 2500 switch. However, if the switch experiences a reboot, the fast-uplink ports (Mode = **Uplink**) use the longer forwarding delay used by ports on standard 802.1D STP (non fast-uplink). This prevents temporary loops that could otherwise result while the switch is determining the STP status for all ports. That is, on ports configured for fast-uplink STP, the first STP state transition after a reboot takes the same amount of time as for redundant ports that are not configured for fast-uplink STP.

Using Fast Uplink with Port Trunks. To use a port trunk for fast-uplink STP, configure it in the same way that you would an individual port for the same purpose. A port trunk configured for fast uplink operates in the same way as an individual, non-trunked port operates; that is, as a logical port.

Note

When you add a port to a trunk, the port takes on the STP mode configured for the trunk, regardless of which STP mode was configured on the port before it was added to the trunk. Thus, all ports belonging to a trunk configured with **Uplink** in the STP **Mode** field will operate in the fast-uplink mode. (If you remove a port from a trunk, the port reverts to the STP Mode setting it had before you added the port to the trunk.)

To use fast uplink over a trunk, you must:

1. Create the trunk.
2. Configure the trunk for fast uplink in the same way that you would configure an individual port for fast uplink.

When you first create a port trunk, its STP Mode setting will be **Norm**, regardless of whether one or more ports in the trunk are set to fast uplink (Mode = **Uplink**). You must still specifically configure the trunk Mode setting to **Uplink**. Similarly, if you eliminate a trunk, the Mode setting on the individual ports in the trunk will return to their previous settings.

Fast-Uplink Troubleshooting

Some of the problems that can result from incorrect useage of Fast-Uplink STP include temporary loops and generation of duplicate packets.

Problem sources can include:

- Fast-Uplink is configured on a switch that is the STP 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 Series 2500 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 STP 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 STP root switch has ports configured for fast-uplink STP and has become the root device due to a failure in the original root device.

The Show Tech Command for Listing Switch Configuration and Operating Details

The **show tech** command provides a tool for gathering information to help with troubleshooting. This command outputs, in a single listing, switch operating and running configuration details 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)
- Stacking status — this switch
- Stacking status — all

Syntax: show tech

Executing **show tech** outputs a data listing to your terminal emulator. However, using your terminal emulator's text capture features, you can also save **show tech** data to a text file for viewing, printing, or sending to an associate. 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.)

To Copy show tech output to a Text File. This example uses the Microsoft Windows terminal emulator. To use another terminal emulator application, refer to the documentation provided with that application.

1. In Hyperterminal, click on **Transfer | Capture Text...**

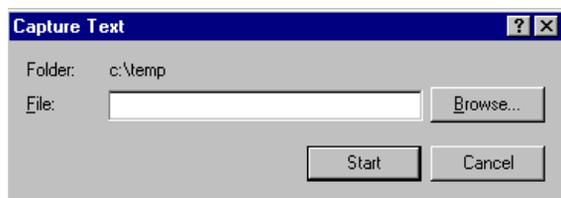


Figure 74. The Capture Text window of the Hypertext Application Used with Microsoft Windows Software

2. In the **File** field, enter the path and file name under which you want to store the **show tech** output.

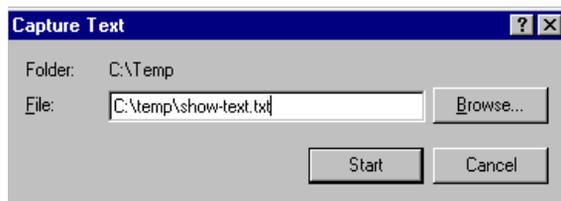


Figure 75. Example of a Path and Filename for Creating a Text File from show tech Output

3. Click **[Start]** to create and open the text file.
4. Execute **show tech**:
HP2512# show tech
 - a. Each time the resulting listing halts and displays `-- MORE --`, press the Space bar to resume the listing.
 - b. When the CLI prompt appears, the show tech listing is complete. At this point, click on **Transfer | Capture Text | Stop** in HyperTerminal to stop copying data into the text file created in the preceding steps.

Note

Remember to do the above step to stop HyperTerminal from copying into the text file. Otherwise, the text file remains open to receiving additional data from the HyperTerminal screen.

5. To access the file, open it in Microsoft Word, Notepad, or a similar text editor.

Enhancements in Release F.02.02

Documentation for Enhancements in Release F.02.02

Software release F.02.02 contains these enhancements:

Enhancement	Summary	Page
TACACS+	TACACS+ authentication enables you to use a central server to allow or deny access to Series 2500 switches (and other TACACS-aware devices) in your network. This means that you can use a central database to create multiple unique username/password sets with associated privilege levels for use by individuals who have reason to access the switch from either the switch's console port (local access) or Telnet (remote access).	154
CDP	In the Series 2500 switches, CDP-v1 (Cisco®Discovery Protocol, version 1) provides data that aids SNMP-based network mapping utilities designed to discover devices running CDP in a network. To make this data available, the switch transmits information about itself via CDP packets to adjacent devices, and also receives and stores information about adjacent devices running CDP. This enables each CDP device to receive and maintain identity data on each of its CDP neighbors and pass this information off to an SNMP utility designed to query the CDP area of the device's MIB.	177
TimeP Change	Changes how to select the TimeP time protocol option.	192
SNTP Time Protocol Enhancement	Adds SNTP, which uses two time protocol operating modes: <ul style="list-style-type: none"> • Broadcast Mode: The switch acquires time updates by accepting the time value from the first SNTP time broadcast detected. • Unicast Mode: The switch requests a time update from the configured SNTP server. 	193
IGMP Enhancements	IGMP on the Series 2500 switches now supports IGMP without IP addressing and Forced Fast-Leave IGMP.	212
Switch Memory Operation	A brief description of how the switch saves configuration changes and a recommendation on when to use the write memory command.	220
Port Security Enhancement	Changes how the switch retains learned static addresses across a reboot.	220
Using the CLI To Configure Usernames	Prior to release F.02.02, you could configure Manager and Operator usernames only from the web browser interface. Beginning with F.02.02 you can also use the CLI to configure usernames.	222

TACACS+ Authentication for Centralized Control of Switch Access Security

TACACS+ Features

Feature	Default	Menu	CLI	Web
view the switch's authentication configuration	n/a	—	page 162	—
view the switch's TACACS+ server contact configuration	n/a	—	page 162	—
configure the switch's authentication methods	disabled	—	page 163	—
configure the switch to contact TACACS+ server(s)	disabled	—	page 166	—

TACACS+ authentication enables you to use a central server to allow or deny access to Series 2500 switches (and other TACACS-aware devices) in your network. This means that you can use a central database to create multiple unique username/password sets with associated privilege levels for use by individuals who have reason to access the switch from either the switch's console port (local access) or Telnet (remote access).

Note

In release F.02.02, TACACS+ authentication does not affect web browser interface access. For steps to block unauthorized access through the web browser interface, see “Controlling Web Browser Interface Access When Using TACACS+ Authentication” on page 173.

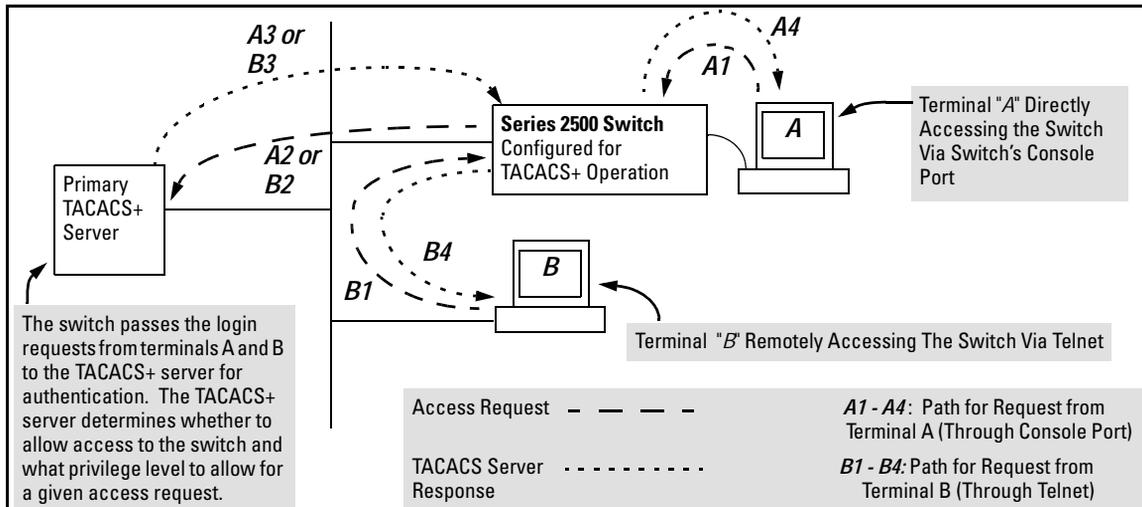


Figure 76. Example of TACACS+ Operation

With authentication configured on the switch and TACACS+ configured and operating on a server in your network, an attempt to log on through Telnet or the switch's serial port will be passed to the TACACS+ server for verification before permission is granted. Similarly, if an operator is using read-only access to the switch and requests read-write access through the CLI **enable** command by entering a user name and password, the switch grants read-write access only after the TACACS+ server verifies the request and returns permission to the switch.

Note

Software release F.02.02 for the Series 2500 switches enables TACACS+ authentication, which is the ability to allow or deny access to a Series 2500 switch on the basis of correct username/password pairs, and to specify the privilege level to allow if access is granted. This release does not support TACACS+ authorization or accounting services.

Series 2500 Switch Authentication Options

With software release F.02.02 installed, the Series 2500 switches include these types of authentication:

- **Local:** Employs a username/password pair assigned locally to the switch. This option allows one username/password pair for manager-level privileges and another username/password pair for operator-level privileges. Local authentication is automatically available in the switch. The *Management and Configuration Guide* you received with your switch describes this method.
- **TACACS+:** Employs a username/password pair assigned remotely to a TACACS+ server application. This option allows multiple username/password pairs for any privilege level available on the switch. The remainder of this section describes TACACS+ authentication on the Series 2500 switches.
- **None:** The switch can be accessed by anyone without requiring a username/password pair. This is the case when TACACS+ is not enabled on the switch and a local, *manager-level* password is not configured in the switch. Allowing the switch to operate in this mode is not recommended because it compromises switch and network access security.

TACACS+ on the Series 2500 switches uses an authentication hierarchy consisting of remote control through a TACACS+ server and the local control (password and user name) built into the switch. That is, with TACACS+ configured on the switch, if the switch cannot contact any designated TACACS+ server, then it defaults to its own locally assigned username/password pairs to control access. To use TACACS+ authentication in a Series 2500 switch, you must enable TACACS+ in the switch and also purchase, install, and configure a third-party TACACS+ server application on the device(s) in your network that you want to use for managing TACACS+ authentication.

Terminology Used in TACACS Applications:

- **NAS (Network Access Server):** This is an industry term for a TACACS-aware device that communicates with a TACACS server for authentication services. Some other terms you may see in literature describing TACACS operation are *communication server*, *remote access server*, or *terminal server*. These terms apply to a Series 2500 switch when TACACS+ is enabled on the switch (that is, when the switch is TACACS-aware).
- **TACACS+ Server:** The server or management station configured as an access control server for TACACS-enabled devices. To use TACACS+ with the Series 2500 switches and any other TACACS-capable devices in your network, you must purchase, install, and configure a TACACS+ server application on a networked server or management station in the network. The TACACS+ server application you install will provide various options for access control and access notifications. For more on the TACACS+ services available to you, see the documentation provided with the TACACS+ server application you will use.
- **Authentication:** The process for granting user access to a device through entry of a user name and password and comparison of this username/password pair with previously stored username/password data. Authentication also grants levels of access, depending on the privileges assigned to a user name and password pair by a system administrator.
 - **Local Authentication:** This method uses username/password pairs configured locally on the switch; one pair each for manager-level and operator-level access to the switch. You can assign local usernames and passwords through the CLI or web browser interface. (Using the menu interface you can assign a local password, but not a user-name.) Because this method assigns passwords to the switch instead of to individuals who access the switch, you must distribute the password information on each switch to everyone who needs to access the switch, and you must configure and manage password protection on a per-switch basis. (For more on local authentication, see the password and username information in the *Configuration and Management Guide* shipped with your Series 2500 switch.
 - **TACACS+ Authentication:** This method enables you to use a TACACS+ server in your network to assign a unique password, user name, and privilege level to each individual or group who needs access to one or more switches or other TACACS-aware devices. This allows you to administer primary authentication from a central server, and to do so with more options than you have when using only local authentication. (You will still need to use local authentication as a backup if your TACACS+ servers become unavailable.) This means, for example, that you can use a central TACACS+ server to grant, change, or deny access to a specific individual on a specific switch instead of having to change local user name and password assignments on the switch itself, and then have to notify other users of the change.

General System Requirements

To use TACACS+ authentication, you need the following:

- Release F.02.02 or later software running on your Series 2500 switch. Ensure that software release F.02.02 or later is running on your switch. Use any of the following methods to view the current software version:

CLI:

```
HP2512> show version
```

Menu Interface:

From the Main Menu, click on

1. **Status and Counters . . .**
 1. **General System Information**

(Check the version number on the **Firmware revision** line.)

Web Browser Interface:

Click on the **Identity** tab.

- A TACACS+ server application installed and configured on one or more servers or management stations in your network. (There are several TACACS+ software packages available.)
- A switch configured for TACACS+ authentication, with access to one or more TACACS+ servers.

Note

The Series 2500 switches include the capability of configuring multiple backup TACACS+ servers. HP recommends that you use a TACACS+ server application that supports a redundant backup installation. This allows you to configure the switch to use a backup TACACS+ server if it loses access to the first-choice TACACS+ server.

TACACS+ Operation

TACACS+ in Series 2500 switches manages authentication of logon attempts through either the Console port or Telnet. For both Console and Telnet you can configure a login (read-only) and an enable (read/write) privilege level access. When your primary authentication control for switch access is a TACACS+ server, you can also specify a local (switch-based) secondary authentication control.

Note

In release F.02.02, TACACS+ does not affect web browser interface access. See "Controlling Web Browser Interface Access" on page 173.

General Authentication Setup Procedure

It is important to test the TACACS+ service before fully implementing it. Depending on the process and parameter settings you use to set up and test TACACS+ authentication in your network, you could accidentally lock all users, including yourself, out of access to a switch. While recovery is simple, it may pose an inconvenience that can be avoided. To prevent an unintentional lockout on a Series 2500 switch, use a procedure that configures and tests TACACS+ protection for one access type (for example, Telnet access), while keeping the other access type (console, in this case) open in case the Telnet access fails due to a configuration problem. The following procedure outlines a general setup procedure.

Note

If a complete access lockout occurs on the switch as a result of a TACACS+ configuration, see "Troubleshooting TACACS+ Operation" on page 175 for recovery methods.

1. Familiarize yourself with the requirements for configuring your TACACS+ server application to respond to requests from a Series 2500 switch. (Refer to the documentation provided with the TACACS+ server software.) This includes knowing whether you need to configure an encryption key. (See "Using the Encryption Key" on page 172.)

2. Ensure that the switch is configured to operate on your network and can communicate with your first-choice TACACS+ server. (At a minimum, this requires IP addressing and a successful **ping** test from the switch to the server.)
3. Determine the following:
 - The IP address(es) of the TACACS+ server(s) you want the switch to use for authentication. If you will use more than one server, determine which server is your first-choice for authentication services.
 - The encryption key, if any, that should be used to allow the switch to communicate with the server.
 - The period you want the switch to wait for a reply to an authentication request before trying another server.
 - The username/password pairs you want the TACACS+ server to use for controlling access to the switch.
 - The privilege level you want for each username/password pair administered by the TACACS+ server for controlling access to the switch.
 - The username/password pairs you want to use for local authentication (one pair each for Operator and Manager levels).
4. Plan and enter the TACACS+ server configuration needed to support TACACS+ operation for Telnet access (login and enable) to the switch. This includes the username/password sets for logging in at the Operator (read-only) privilege level and the sets for logging in at the Manager (read/write) privilege level.

Note on Privilege Levels

When a TACACS+ server authenticates an access request from a switch, it includes a privilege level code for the switch to use in determining which privilege level to grant to the terminal requesting access. The switch interprets a privilege level code of "15" as authorization for the Manager (read/write) privilege level access. Privilege level codes of 14 and lower result in Operator (read-only) access. Thus, when configuring the TACACS+ server response to a request that includes a username/password pair that should have Manager privileges, you must use a privilege level of 15. For more on this topic, refer to the documentation you received with your TACACS+ server application.

If you are a first-time user of the TACACS+ service, HP recommends that you configure only the minimum feature set required by the TACACS+ application to provide service in your network environment. After you have success with the minimum feature set, you may then want to try additional features that the application offers.

5. Ensure that the switch has the correct local username and password for Manager access. (If the switch cannot find any designated TACACS+ servers, the local manager and operator username/password pairs are always used as the secondary access control method.)

Caution

You should ensure that the switch has a local Manager password. Otherwise, if authentication through a TACACS+ server fails for any reason, then unauthorized access will be available through the console port or Telnet.

6. Using a terminal device connected to the switch's console port, configure the switch for TACACS+ authentication *only* for **telnet login** access and **telnet enable** access. At this stage, do not configure TACACS+ authentication for console access to the switch, as you may need to use the console for access if the configuration for the Telnet method needs debugging.
7. On a remote terminal device, use Telnet to attempt to access the switch. If the attempt fails, use the console access to check the TACACS+ configuration on the switch. If you make changes in the switch configuration, check Telnet access again. If Telnet access still fails, check the configuration in your TACACS+ server application for mis-configurations or missing data that could affect the server's interoperation with the switch.
8. After your testing shows that Telnet access using the TACACS+ server is working properly, configure your TACACS+ server application for console access. Then test the console access. If access problems occur, check for and correct any problems in the switch configuration, and then test console access again. If problems persist, check your TACACS+ server application for mis-configurations or missing data that could affect the console access.
9. When you are confident that TACACS+ access through both Telnet and the switch's console operates properly, use the **write memory** command to save the switch's running-config file to flash.

Configuring TACACS+ on the Switch

The switch offers three command areas for TACACS+ operation:

- **show authentication** and **show tacacs**: Displays the switch's TACACS+ configuration and status.
- **aaa authentication**: A command for configuring the switch's authentication methods
- **tacacs-server**: A command for configuring the switch's contact with TACACS+ servers

CLI Commands Described in this Section

show authentication	below
show tacacs	page 162
aaa authentication	pages 163 through 165
console	pages 163 through 165
Telnet	pages 163 through 165
num-attempts <1..10>	pages 163 through 165
tacacs-server	pages 166 through 168
host <ip addr>	pages 166 through 168
key	page 168
timeout <1 ..255>	page 169

Viewing the Switch's Current Authentication Configuration

This command lists the number of login attempts the switch allows in a single login session, and the primary/secondary access methods configured for each type of access.

Syntax: show authentication

This example shows the default authentication configuration.

```

HP2512> show authentication
Status and Counters - Authentication Information
Login Attempts : 3

      Login   Login   Enable  Enable
Access Task Primary Secondary Primary Secondary
-----
Console Local   None    Local   None
Telnet  Local   None    Local   None
    
```

Configuration for login and enable access to the switch through the switch console port.

Configuration for login and enable access to the switch through Telnet.

Figure 77. Example Listing of the Switch's Authentication Configuration

Viewing the Switch's Current TACACS+ Server Contact Configuration

This command lists the timeout period, encryption key, and the IP addresses of the first-choice and backup TACACS+ servers the switch can contact.

Syntax: show tacacs

For example, if the switch was configured for a first-choice and two backup TACACS+ server addresses, the default timeout period, and **paris-1** for a (global) encryption key, **show tacacs** would produce a listing similar to the following:

```

HP2512 (config)# show tacacs
Status and Counters - TACACS Information
Timeout : 5
Encryption Key : paris-1
Server IP Addr  Opens  Closes  Aborts  Errors  Pkts Rx  Pkts Tx
-----
10.30.248.100  0      0      0      0      0      0
10.30.248.156  0      0      0      0      0      0
10.30.248.105  0      0      0      0      0      0
    
```

First-Choice TACACS+ Server

Second-Choice TACACS+ Server

Third-Choice TACACS+ Server

Figure 78. Example of the Switch's TACACS+ Configuration Listing

Configuring the Switch's Authentication Methods

The **aaa authentication** command configures the access control for console port and Telnet access to the switch. That is, for both access methods, aaa authentication specifies whether to use a TACACS+ server or the switch's local authentication, or (for some secondary scenarios) no authentication (meaning that if the primary method fails, authentication is denied). This command also reconfigures the number of access attempts to allow in a session if the first attempt uses an incorrect username/password pair.

Syntax: aaa authentication < console | telnet > < enable | login > < local | tacacs > < local | none >
 aaa authentication num-attempts < 1..10 >

Table 12. AAA Authentication Parameters

Name	Default	Range	Function
console - or - telnet	n/a	n/a	Specifies whether the command is configuring authentication for the console port or Telnet access method for the switch.
enable - or - login	n/a	n/a	Specifies the privilege level for the access method being configured. login: Operator (read-only) privileges enable: Manager (read-write) privileges
local - or - tacacs	local	n/a	Specifies the primary method of authentication for the access method being configured. local: Use the username/password pair configured locally in the switch for the privilege level being configured tacacs: Use a TACACS+ server.
local - or - none	none	n/a	Specifies the secondary (backup) type of authentication being configured. local: The username/password pair configured locally in the switch for the privilege level being configured none: No secondary type of authentication for the specified method/privilege path. (Available only if the primary method of authentication for the access being configured is local.) Note: If you do not specify this parameter in the command line, the switch automatically assigns the secondary method as follows: <ul style="list-style-type: none"> • If the primary method is tacacs, the only secondary method is local. • If the primary method is local, the default secondary method is none.
num-attempts	3	1 - 10	In a given session, specifies how many tries at entering the correct username/password pair are allowed before access is denied and the session terminated.

As shown in the following table, login and enable access is always available locally through a direct terminal connection to the switch's console port. However, for Telnet access, you can configure TACACS+ to deny access if a TACACS+ server goes down or otherwise becomes unavailable to the switch.

Table 13. Primary/Secondary Authentication Table

Access Method and Privilege Level	Authentication Options		Effect on Access Attempts
	Primary	Secondary	
Console — Login	local	none*	Local username/password access only.
	tacacs	local	If Tacacs+ server unavailable, uses local username/password access.
Console — Enable	local	none*	Local username/password access only.
	tacacs	local	If Tacacs+ server unavailable, uses local username/password access.
Telnet — Login	local	none*	Local username/password access only.
	tacacs	local	If Tacacs+ server unavailable, uses local username/password access.
	tacacs	none	If Tacacs+ server unavailable, denies access.
Telnet — Enable	local	none*	Local username/password access only.
	tacacs	local	If Tacacs+ server unavailable, uses local username/password access.
	tacacs	none	If Tacacs+ server unavailable, denies access.

*When "local" is the primary option, you can also select "local" as the secondary option. However, in this case, a secondary "local" is meaningless because the switch has only one local level of username/password protection.

For example, here is a set of access options and the corresponding commands to configure them:

Console Login (Operator, or Read-Only) Access: Primary using TACACS+ server. Secondary using Local.				
HP2512(config)# aaa authentication console login tacacs local				
	<i>Console Login (Operator, or Read- Only Access)</i>	<i>Primary</i>	<i>Secondary</i>	

Console Enable (Manager, or Read/Write) Access: Primary using TACACS+ server. Secondary using Local.				
HP2512(config)# aaa authentication console enable tacacs local				
	<i>Console Enable (Manager, or Read/ Write Access)</i>	<i>Primary</i>	<i>Secondary</i>	

Telnet Login (Operator, or Read-Only) Access: Primary using TACACS+ server. Secondary using Local.				
HP2512(config)# aaa authentication telnet login tacacs local				
	<i>Telnet Login (Operator, or Read- Only Access)</i>	<i>Primary</i>	<i>Secondary</i>	

Telnet Enable (Manager, or Read/Write) Access: Primary using TACACS+ server. Secondary using Local.				
HP2512(config)# aaa authentication telnet enable tacacs local				
	<i>Telnet Enable (Manager, or Read/ Write Access)</i>	<i>Primary</i>	<i>Secondary</i>	

Deny Access and Close the Session After Failure of Two Consecutive Username/Password Pairs:				
HP2512(config)# aaa authentication num-attempts 2				
	<i>Attempt Limit</i>			

Configuring the Switch's TACACS+ Server Access

The `tacacs-server` command configures these parameters:

- **The host IP address(es)** for up to three TACACS+ servers; one first-choice and up to two backups. Designating backup servers provides for a continuation of authentication services in case the switch is unable to contact the first-choice server.
- **An optional encryption key.** This key helps to improve security, and must match the encryption key used in your TACACS+ server application. In some applications, the term "secret key" or "secret" may be used instead of "encryption key". If you need only one encryption key for the switch to use in all attempts to authenticate through a TACACS+ server, configure a global key. However, if the switch is configured to access multiple TACACS+ servers having different encryption keys, you can configure the switch to use different encryption keys for different TACACS+ servers.
- **The timeout value** in seconds for attempts to contact a TACACS+ server. If the switch sends an authentication request, but does not receive a response within the period specified by the timeout value, the switch resends the request to the next server in its Server IP Addr list, if any. If the switch still fails to receive a response from any TACACS+ server, it reverts to whatever secondary authentication method was configured using the **aaa authentication** command (local or none; see “Configuring the Switch's Authentication Methods” on page 163.)

Syntax:	<code>tacacs-server host <ip-addr> [key <key-string>]</code>	<i>Adds a TACACS+ server and optionally assigns a server-specific encryption key.</i>
	<code>[no] tacacs-server host <ip-addr></code>	<i>Removes a TACACS+ server assignment (including its encryption key, if any).</i>
	<code>tacacs-server key <key-string></code>	<i>Enters the optional global encryption key.</i>
	<code>[no] tacacs-server key</code>	<i>Removes the optional global encryption key. (Does not affect any server-specific encryption key assignments.)</i>
	<code>tacacs-server timeout < 1 . . 255 ></code>	<i>Changes the wait period for a TACACS server response. (Default: 5 seconds.)</i>

Name	Default	Range
host <ip-addr> [key <key-string>	none	n/a

Specifies the IP address of a device running a TACACS+ server application. Optionally, can also specify the unique, per-server encryption key to use when each assigned server has its own, unique key. For more on the encryption key, see "Using the Encryption Key" on page 172 and the documentation provided with your TACACS+ server application.

You can enter up to three IP addresses; one first-choice and two (optional) backups (second-choice and third-choice). Use **show tacacs** to view the current IP address list.

If the first-choice TACACS+ server fails to respond to a request, the switch tries the second address, if any, in the show tacacs list. If the second address also fails, then the switch tries the third address, if any. (See figure 78, "Example of the Switch's TACACS+ Configuration Listing" on page 162.)

The priority (first-choice, second-choice, and third-choice) of a TACACS+ server in the switch's TACACS+ configuration depends on the order in which you enter the server IP addresses:

1. When there are no TACACS+ servers configured, entering a server IP address makes that server the first-choice TACACS+ server.
 2. When there is one TACACS+ server already configured, entering another server IP address makes that server the second-choice (backup) TACACS+ server.
 3. When there are two TACACS+ servers already configured, entering another server IP address makes that server the third-choice (backup) TACACS+ server.
- The above position assignments are fixed. Thus, if you remove one server and replace it with another, the new server assumes the priority position that the removed server had. For example, suppose you configured three servers, A, B, and C, configured in order:
First-Choice:A; Second-Choice:B; Third-Choice: C
 - If you removed server B and then entered server X, the TACACS+ server order of priority would be:
First-Choice:A; Second-Choice:X; Third-Choice: C
 - If there are two or more vacant slots in the TACACS+ server priority list and you enter a new IP address, the new address will take the vacant slot with the highest priority. Thus, if A, B, and C are configured as above and you (1) remove A and B, and (2) enter X and Y (in that order), then the new TACACS+ server priority list would be X, Y, and C.
 - The easiest way to change the order of the TACACS+ servers in the priority list is to remove all server addresses in the list and then re-enter them in order, with the new first-choice server address first, and so on.

To add a new address to the list when there are already three addresses present, you must first remove one of the currently listed addresses.

See also "General Authentication Process Using a TACACS+ Server" on page 170.

key < key-string >	none (null)	n/a
---------------------------	-------------	-----

Specifies the optional, global "encryption key" that is also assigned in the TACACS+ server(s) that the switch will access for authentication. This option is subordinate to any "per-server" encryption keys you assign, and applies only to accessing TACACS+ servers for which you have not given the switch a "per-server" key. (See the **host** <ip-addr> [**key** <key-string> entry at the beginning of this table.)

For more on the encryption key, see "Using the Encryption Key" on page 172 and the documentation provided with your TACACS+ server application.

timeout < 1..255 >	5 sec	1 - 255 sec
---------------------------	-------	-------------

Specifies how long the switch waits for a TACACS+ server to respond to an authentication request. If the switch does not detect a response within the timeout period, it initiates a new request to the next TACACS+ server in the list. If all TACACS+ servers in the list fail to respond within the timeout period, the switch uses either local authentication (if configured) or denies access (if **none** configured for local authentication).

Adding, Removing, or Changing the Priority of a TACACS+ Server. Suppose that the switch was already configured to use TACACS+ servers at 10.28.227.10 and 10.28.227.15. In this case, 10.28.227.15 was entered first, and so is listed as the first-choice server:

```
HP2512 (config)# show tacacs
Status and Counters - TACACS Information
Timeout : 5
Encryption Key : First-Choice TACACS+ Server
```

Server IP Addr	Closes	Aborts	Errors	Pkts Rx	Pkts Tx
10.28.227.15	0	0	0	0	0
10.28.227.10	0	0	0	0	0

Figure 79. Example of the Switch with Two TACACS+ Server Addresses Configured

To move the "first-choice" status from the "15" server to the "10" server, use the **no tacacs-server host <ip-addr>** command to delete both servers, then use **tacacs-server host <ip-addr>** to re-enter the "10" server first, then the "15" server.

The servers would then be listed with the new "first-choice" server, that is:

```
HP2512 (config)# show tacacs
Status and Counters - TACACS Information
Timeout : 5
Encryption Key :
```

Server IP Addr	Opens	Closes	Aborts	Errors	Pkts Rx	Pkts Tx
10.28.227.10	0	0	0	0	0	0
10.28.227.15	0	0	0	0	0	0

The "10" server is now the "first-choice" TACACS+ authentication device.

Figure 80. Example of the Switch After Assigning a Different "First-Choice" Server

To remove the 10.28.227.15 device as a TACACS+ server, you would use this command:

```
HP2512 (config)# no tacacs-server host 10.28.227.15
```

Configuring an Encryption Key. Use an encryption key in the switch if the switch will be requesting authentication from a TACACS+ server that also uses an encryption key. (If the server expects a key, but the switch either does not provide one, or provides an incorrect key, then the authentication attempt will fail.) Use a *global encryption key* if the same key applies to all TACACS+ servers the switch may use for authentication attempts. Use a *per-server encryption key* if different servers the switch may use will have different keys. (For more details on encryption keys, see "Using the Encryption Key" on page 172.)

To configure **westside** as a global encryption key:

```
HP2512(config) tacacs-server key westside
```

To configure **westside** as a per-server encryption key:

```
HP2512(config) tacacs-server host 10.28.227.63 key westside
```

An encryption key can contain up to 100 characters, without spaces, and is likely to be case-sensitive in most TACACS+ server applications.

To delete a global encryption key from the switch, use this command:

```
HP2512(config)# no tacacs-server key
```

To delete a per-server encryption key in the switch, re-enter the `tacacs-server host` command without the key parameter. For example, if you have **westside** configured as the encryption key for a TACACS+ server with the IP address of 10.28.227.104 and you wanted to eliminate the key, you would use this command:

```
HP2512(config)# tacacs-server host 10.28.227.104
```

Note

The `show tacacs` command lists the global encryption key, if configured. However, to view any configured per-server encryption keys, you must use **show config running**.

Configuring the Timeout Period. The timeout period specifies how long the switch waits for a response to an authentication request from a TACACS+ server before either sending a new authentication request to the next server in the switch's Server IP Address list or using the local authentication option. For example, to change the timeout period from 5 seconds (the default) to 3 seconds:

```
HP2512(config)# tacacs-server timeout 3
```

How Authentication Operates

General Authentication Process Using a TACACS+ Server

Authentication through a TACACS+ server operates generally as described below. For specific operating details, refer to the documentation you received with your TACACS+ server application.

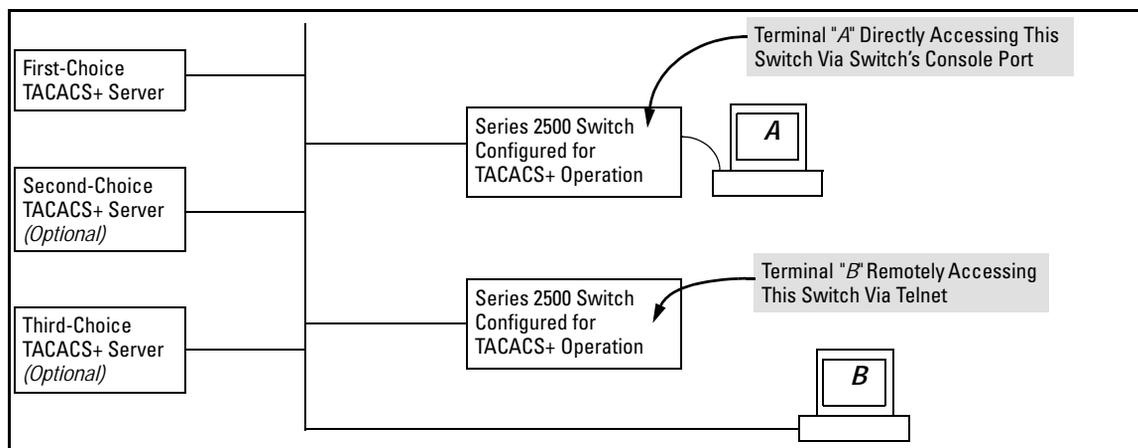


Figure 81. Using a TACACS+ Server for Authentication

Using figure 81, above, after either switch detects an operator's logon request from a remote or directly connected terminal, the following events occur:

1. The switch queries the first-choice TACACS+ server for authentication of the request.
 - If the switch does not receive a response from the first-choice TACACS+ server, it attempts to query a secondary server. If the switch does not receive a response from any TACACS+ server, then it uses its own local username/password pairs to authenticate the logon request. (See "Local Authentication Process", on page 171.)
 - If a TACACS+ server recognizes the switch, it forwards a username prompt to the requesting terminal via the switch.
2. When the requesting terminal responds to the prompt with a username, the switch forwards it to the TACACS+ server.
3. After the server receives the username input, the requesting terminal receives a password prompt from the server via the switch.
4. When the requesting terminal responds to the prompt with a password, the switch forwards it to the TACACS+ server and one of the following actions occurs:

- If the username/password pair received from the requesting terminal matches a username/password pair previously stored in the server, then the server passes access permission through the switch to the terminal.
- If the username/password pair entered at the requesting terminal does not match a username/password pair previously stored in the server, access is denied. In this case, the terminal is again prompted to enter a username and repeat steps 2 through 4. In the default configuration, the switch allows up to three attempts to authenticate a login session. If the requesting terminal exhausts the attempt limit without a successful TACACS+ authentication, the login session is terminated and the operator at the requesting terminal must initiate a new session before trying again.

Local Authentication Process

When the switch is configured to use TACACS+, it reverts to local authentication only if one of these two conditions exists:

- "Local" is the authentication option for the access method being used.
- The switch has been configured to query a TACACS+ server for an authentication request, but has not received a response

(For a listing of authentication options, see Table 13 on page 164.)

For local authentication, the switch uses the operator-level and manager-level username/password set(s) previously configured locally on the switch. (These are the usernames and passwords you can configure using the CLI password command, the web browser interface, or the menu interface—which enables only local password configuration).

- If the operator at the requesting terminal correctly enters the username/password pair for either access level, access is granted.
- If the username/password pair entered at the requesting terminal does not match either username/password pair previously configured locally in the switch, access is denied. In this case, the terminal is again prompted to enter a username/password pair. In the default configuration, the switch allows up to three attempts. If the requesting terminal exhausts the attempt limit without a successful authentication, the login session is terminated and the operator at the requesting terminal must initiate a new session before trying again.

Using the Encryption Key

General Operation

When used, the encryption key (sometimes termed "key", "secret key", or "secret") helps to prevent unauthorized intruders on the network from reading username and password information in TACACS+ packets moving between the switch and a TACACS+ server. At the TACACS+ server, a key may include both of the following:

- **Global key:** A general key assignment in the TACACS+ server application that applies to all TACACS-aware devices for which an individual key has not been configured.
- **Individual key:** A unique key assignment in the TACACS+ server application that applies to a specific TACACS-aware device.

Note

Configure a key in the switch only if the TACACS+ server application has this exact same key configured for the switch. That is, if the key parameter in switch "X" does not exactly match the key setting for switch "X" in the TACACS+ server application, then communication between the switch and the TACACS+ server will fail and authentication results will be unpredictable.

Thus, on the TACACS+ server side, you have a choice as to how to implement a key. On the switch side, it is necessary only to enter the key parameter so that it exactly matches its counterpart in the server. For information on how to configure a general or individual key in the TACACS+ server, refer to the documentation you received with the application.

Encryption Options in the Switch

When configured, the encryption key causes the switch to encrypt the TACACS+ packets it sends to the server. When left at "null", the TACACS+ packets are sent in clear text. The encryption key (or just "key") you configure in the switch must be identical to the encryption key configured in the corresponding TACACS+ server. If the key is the same for all TACACS+ servers the switch will use for authentication, then configure a global key in the switch. If the key is different for one or more of these servers, use "per-server" keys in the switch. (If you configure both a global key and one or more per-server keys, the per-server keys will override the global key for the specified servers.)

For example, you would use the next command to configure a global encryption key in the switch to match a key entered as **north40campus** in two target TACACS+ servers. (That is, both servers use the same key for your switch.) Note that you do not need the server IP addresses to configure a global key in the switch:

```
HP2512(config)# tacacs-server key north40campus
```

Suppose that you subsequently add a third TACACS+ server (with an IP address of 10.28.227.87) that has **south10campus** for an encryption key. Because this key is different than the one used for the two servers in the previous example, you will need to assign a per-server key in the switch that applies only to the designated server:

```
HP2512(config)# tacacs-server host 10.28.227.87 key south10campus
```

With both of the above keys configured in the switch, the **south10campus** key overrides the **north40campus** key only when the switch tries to access the TACACS+ server having the 10.28.227.87 address.

Controlling Web Browser Interface Access When Using TACACS+ Authentication

In release F.02.02, configuring the switch for TACACS+ authentication does not affect web browser interface access. To prevent unauthorized access through the web browser interface, do one or more of the following:

- Configure local authentication (a Manager user name and password and, optionally, an Operator user name and password) on the switch.
- Configure the switch's Authorized IP Manager feature to allow web browser access only from authorized management stations. (The Authorized IP Manager feature does not interfere with TACACS+ operation.)
- Disable web browser access to the switch.

Messages

The switch generates the CLI messages listed below. However, you may see other messages generated in your TACACS+ server application. For information on such messages, refer to the documentation you received with the application.

Table 14. Tacacs Messages

CLI Message	Meaning
Connecting to Tacacs server	The switch is attempting to contact the TACACS+ server identified in the switch's tacacs-server configuration as the first-choice (or only) TACACS+ server.
Connecting to secondary Tacacs server	The switch was not able to contact the first-choice TACACS+ server, and is now attempting to contact the next (secondary) TACACS+ server identified in the switch's tacacs-server configuration.
Invalid password	The system does not recognize the username or the password or both. Depending on the authentication method (tacacs or local), either the TACACS+ server application did not recognize the username/password pair or the username/password pair did not match the username/password pair configured in the switch.
No Tacacs servers responding	The switch has not been able to contact any designated TACACS+ servers. If this message is followed by the Username prompt, the switch is attempting local authentication.
Not legal combination of authentication methods	For console access, if you select tacacs as the primary authentication method, you must select local as the secondary authentication method. This prevents you from being locked out of the switch if all designated TACACS+ servers are inaccessible to the switch.
Record already exists	When resulting from a tacacs-server host <ip addr> command, indicates an attempt to enter a duplicate TACACS+ server IP address.

Operating Notes

- If you configure Authorized IP Managers on the switch, it is not necessary to include any devices used as TACACS+ servers in the authorized manager list. That is, TACACS+ operates regardless of any Authorized IP Manager configuration.
- When TACACS+ is not enabled on the switch—or when the switch's only designated TACACS+ servers are not accessible— *setting a local Operator password without also setting a local Manager password does not protect the switch from manager-level access by unauthorized persons.*)

Troubleshooting TACACS+ Operation

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 due to 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, then pressing the Reset button or cycling the power reboots the switch with the boot-up configuration.
- Disconnect the switch from network access to any TACACS+ servers and then log in to the switch using either Telnet or direct console port access. Because the switch cannot access a TACACS+ server, it will default 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), then 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 key.)
- 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.

Enhancements in Release F.02.02

- The time quota for the account has been exhausted.
- The time credit for the account has expired.
- The access attempt is outside of the timeframe allowed for the account.
- The allowed number of concurrent logins for the account has been exceeded

For more help, refer to 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. Refer to 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.

CDP

CDP Features

Feature	Default	Menu	CLI	Web
view the switch's CDP configuration	n/a	—	page 182	—
view the switch's CDP Neighbors table	n/a	—	page 183	—
clear (reset) the CDP Neighbors table	n/a	—	page 184	—
enable or disable CDP on the switch	enabled	—	page 185	—
enable or disable CDP operation on an individual port	enabled	—	page 186	—
change the transmit interval for the switch's CDP packets	60 seconds	—	page 186	—
change the hold time (time-to-live for CDP packets the switch generates)	180 seconds	—	page 187	—

Introduction

In the Series 2500 switches, CDP-v1 (Cisco Discovery Protocol, version 1) provides data that aids SNMP-based network mapping utilities designed to discover devices running CDP in a network. To make this data available, the switch transmits information about itself via CDP packets to adjacent devices, and also receives and stores information about adjacent devices running CDP. This enables each CDP device to receive and maintain identity data on each of its CDP neighbors and pass this information off to an SNMP utility designed to query the CDP area of the device's MIB.

Note

To take advantage of CDP in the Series 2500 switches, you should have a working knowledge of SNMP operation and an SNMP utility capable of polling the switches for CDP data. HP's implementation of CDP places specific data into the switch's Management Information Base (MIB). However, retrieval of this data for network mapping is dependent on the operation of your SNMP utility. Refer to the documentation provided with the utility.

An SNMP utility can progressively discover CDP devices in a network by:

1. Reading a given device's CDP Neighbor table (in the Management Information Base, or MIB) to learn about other, neighbor CDP devices
2. Using the information learned in step 1 to go to and read the neighbor devices' CDP Neighbors tables to learn about additional CDP devices, and so on

This section describes CDP operation in the Series 2500 switches. For information on how to use an SNMP utility to retrieve the CDP information from the switch's CDP Neighbors table (in the switch's MIB), refer to the documentation provided with the particular SNMP utility. For information on the object identifiers in the CDP MIB, see "CDP Neighbor Data and MIB Objects" on page 188.

CDP Terminology

- **CDP Device:** A switch, server, router, workstation, or other device running CDP.
- **CDP-Aware:** A device that has CDP in its operating code (with CDP either enabled or disabled in that device).
- **CDP-Disabled:** A CDP-aware device on which CDP is currently disabled.
- **Non-CDP Device:** A device that does not have CDP in its operating code.
- **CDP Neighbor:** A CDP device that is either directly connected to another CDP device or connected to that device by a non-CDP device, such as some hubs.

General CDP Operation

The switch stores information about adjacent CDP devices in a *CDP Neighbors table*. For example:

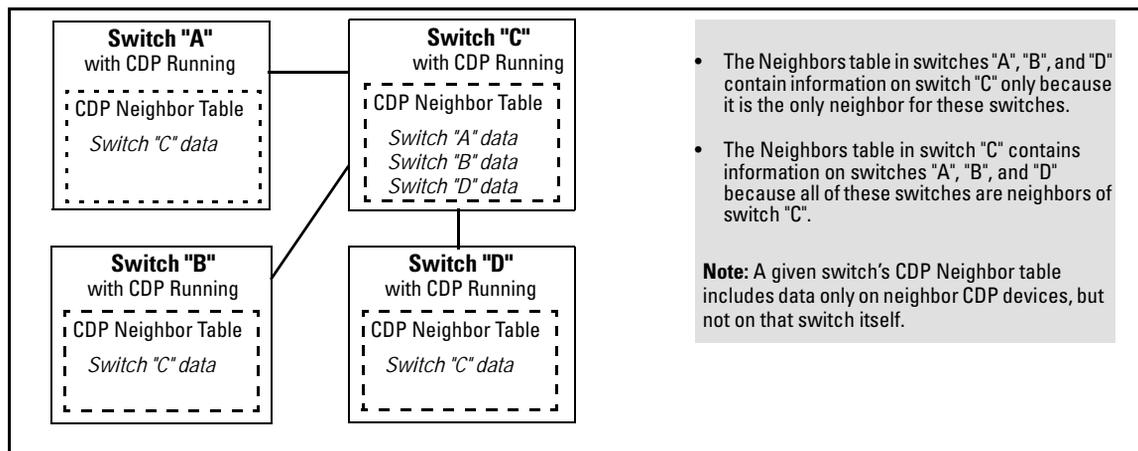


Figure 82. Example of How Series 2500 Switches Store Data on Neighbor CDP Devices

Outgoing Packets

A Series 2500 switch running CDP periodically transmits a one-hop CDP packet out each of its ports. This packet contains data describing the switch and, if the one-hop destination is another device running CDP, the receiving device stores the sending device's data in a CDP Neighbors table. The receiving device also transmits a similar one-hop CDP packet out each of its ports to make itself known to other CDP devices to which it is connected. Thus, each CDP device in the network provides data on itself to the CDP neighbors to which it is directly connected. However, there are instances where a packet is forwarded beyond the immediate neighbor, or simply dropped.

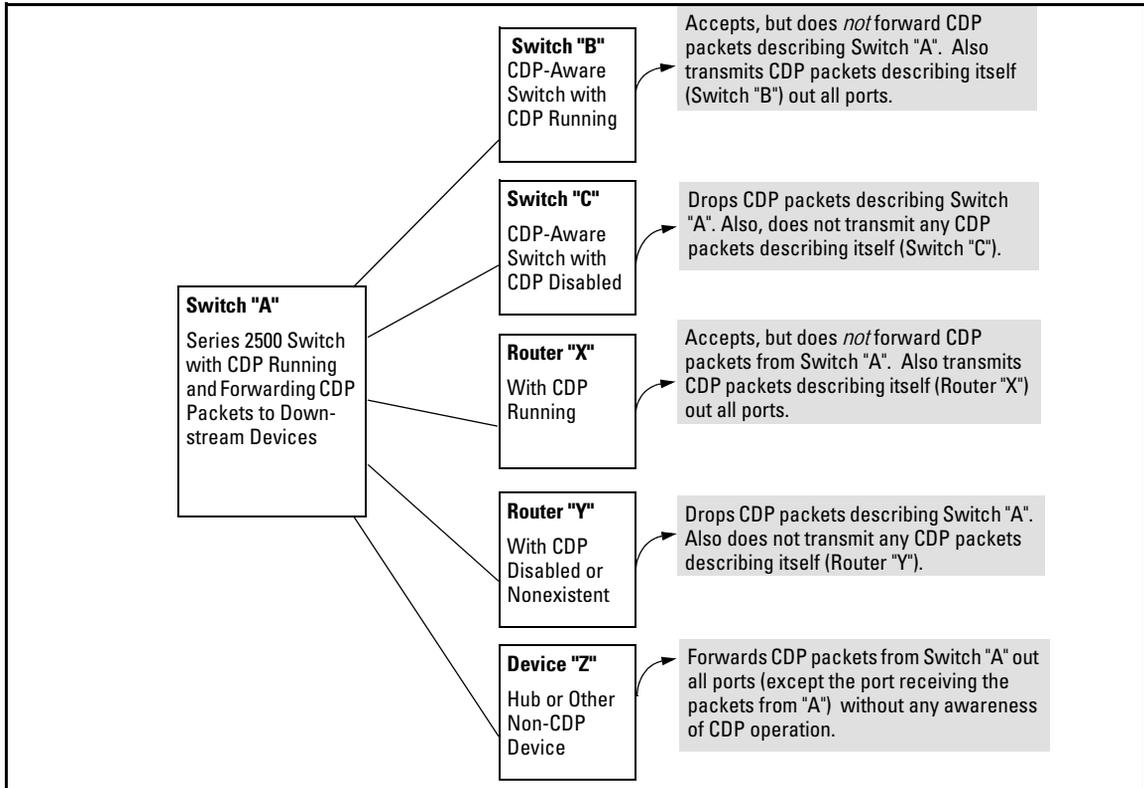


Figure 83. Example of Outgoing CDP Packet Operation

Incoming CDP Packets

When a CDP-enabled Series 2500 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.) The Series 2500 switches purge expired CDP neighbor entries every three seconds.

Non-CDP devices such as some hubs and other devices that do not have CDP capability are transparent to CDP operation. (Other hubs are CDP-aware, but still forward CDP packets as if they were transparent to CDP operation. See "CDP-Capable Hubs" on page 190.) However, an intervening CDP-aware device that is CDP-disabled is *not* transparent. For example, in figure 84, the CDP neighbor pairs are as follows: A/1, A/2, A/3, A/B, B/C. Note that "C" and "E" are *not* neighbors because the intervening CDP-disabled switch "D" does not forward CDP packets; i.e. is not transparent to CDP traffic. (For the same reason, switch "E" does not have any CDP neighbors.)

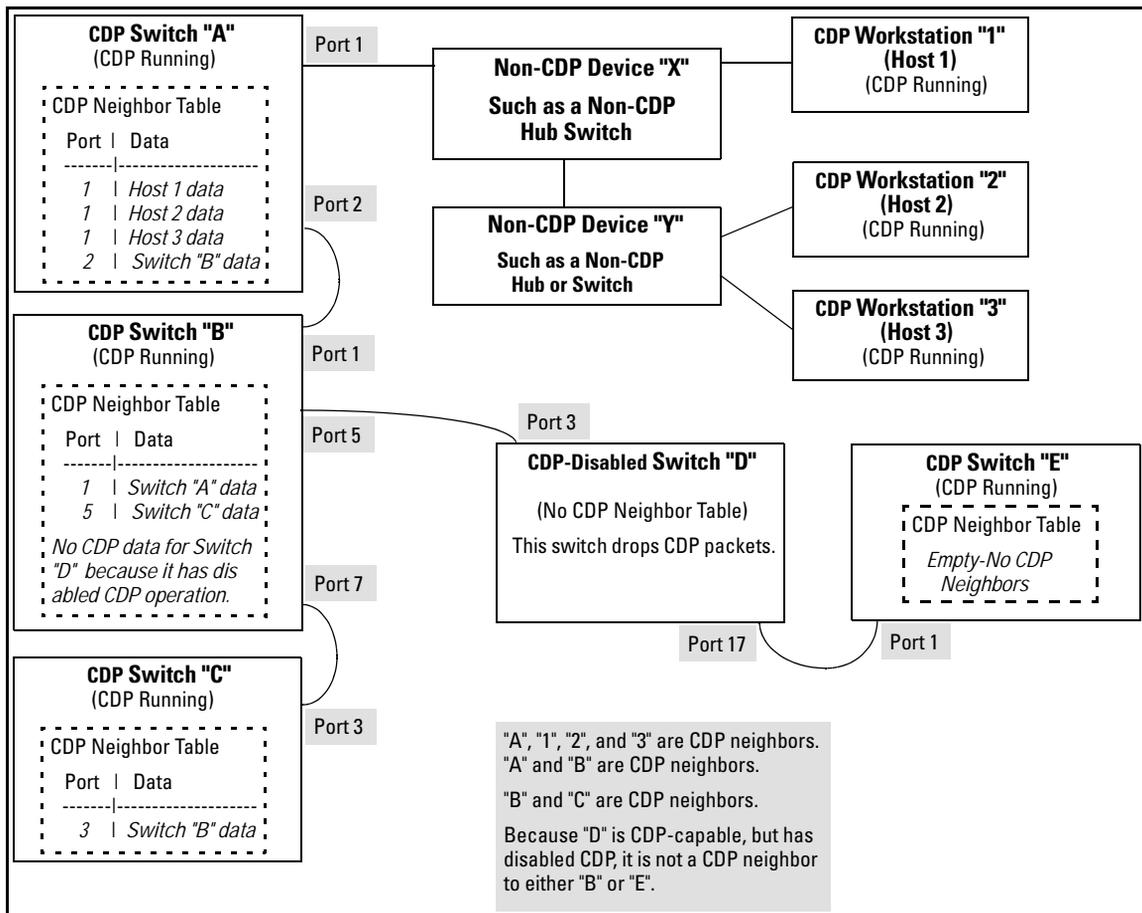


Figure 84. Example of Incoming CDP Packet Results

Using the example in figure 84:

The CDP Neighbor table for switches "A" and "B" would appear similar to these:

Switch A:

Port	Device ID	Platform	Capability
1	XYZ (0050c0-814b01)	XYZ Workstation	H
1	XYZ (0050c0-850a43)	XYZ Workstation	H
1	XYZ (0050c0-850b87)	XYZ Workstation	H
2	HP2512 (0030c1-7fec40)	HP J4812A ProCurve Switch...	S

Switch B:

Port	Device ID	Platform	Capability
1	Switch A (0030c1-583b39)	HP J4812A ProCurve Switch...	S
7	Switch B (0060b0-889e00)	HP J4813A ProCurve Switch...	S

(Note that no CDP devices appear on port 5, which is connected to a device on which CDP is present, but disabled.)

Figure 85. Example of Viewable CDP Neighbor Table for Switches "A" and "B" in Figure 84

Thus, based on the CDP packets it receives, each CDP device maintains a per-port data entry for each of its neighbors that are running CDP, but not for other CDP devices that are accessible only through a CDP neighbor. (See the relationship between switches A, B, and C in figure 84.) In other words, a CDP device will have data on its immediate CDP neighbors (including those reached through a device that is transparent to CDP), but not to other CDP devices in the network.

Table 15. How Devices Handle Incoming CDP Packets

Status of Device Receiving a CDP Packet	Action of Receiving Device
Running CDP	Stores neighbor data in CDP Neighbor table. Does not forward CDP packet.
CDP Disabled	Drops CDP packet. There is no CDP Neighbor table and no CDP neighbor data is stored.
No CDP Capability	Forwards CDP packet out all ports except the port on which the packet was received.
Router Running CDP	Stores neighbor data in CDP Neighbor table. Does not forward CDP packet.
Router with CDP (1) Disabled or (2) Not CDP-Capable	Drops CDP packet.

Non-CDP devices (that is, devices that are not capable of running CDP) are transparent to CDP operation. However, an intervening CDP-aware device that is CDP-disabled is *not* transparent. For example, in figure 84 (page 180), "B", "D", and "E" are *not* CDP neighbors because "D" (the intervening CDP-disabled switch) does not forward CDP packets; i.e. is not transparent to CDP traffic. (For the same reason, switch "E" does not have any CDP neighbors.)

Figure 84 (page 180) illustrates how multiple CDP neighbors can appear on a single port. In this case, switch "A" has three CDP neighbors on port 1 because the intervening devices are not CDP-capable and simply forward CDP neighbors data out all ports (except the port on which the data was received).

Configuring CDP on the Switch

On a Series 2500 switch you can:

- View the switch's current global and per-port CDP configuration
- List the current contents of the switch's CDP Neighbors table (that is, view a listing of the CDP devices of which the switch is aware)
- Enable or disable CDP (Default: Enabled)
- Specify the hold time (CDP packet time-to-live) for CDP data delivered to neighboring CDP devices. For example, in CDP switch "A" you can specify the hold time for switch "A" entries in the CDP Neighbor tables of other CDP devices. (Default: 180 seconds)
- Specify the transmission interval for CDP packets. (Default: 60 seconds)

CLI Commands Described in this Section

show CDP	below
show CDP neighbors	page 183
cdp clear	page 184
[no] cdp run	page 185
[no] cdp enable	page 186
cdp holdtime	page 186
cdp timer	page 187

Viewing the Switch's Current CDP Configuration

This command lists the switch's global and per-port CDP configuration. (In the factory default configuration, the switch runs CDP on all ports with a hold time of 180 seconds and a transmit interval of 60 seconds.)

Syntax: show cdp

This example shows the default CDP configuration.

```

HP2512# show cdp
Enable CDP : Yes
CDP Hold Time : 180
CDP Transmit Interval : 60
Port CDP
-----
1    enabled
2    enabled
3    enabled
.    .
.    .
.    .
14   enabled

```

CDP Enable/Disable on the Switch
 Packet Hold Time in CDP Neighbor Table
 Interval for Transmitting Outbound CDP Packets on All Ports
 Per-Port CDP Enable/Disable

Figure 86. Example of a CDP Configuration Listing

Viewing the Current Contents of the Switch's CDP Neighbors Table

This command lists the neighboring CDP devices the switch has detected. Devices are listed by the port on which they were detected. The entry for a specific device includes a subset of the information collected from the device's CDP packet. (For more on this topic, see “CDP Neighbor Data and MIB Objects” on page 188.)

Syntax: show cdp neighbors

This example lists six CDP devices (four switches and two workstations) that the switch has detected by receiving their CDP packets.

```

HP2512> show cdp neighbors
CDP neighbors information

```

Port	Device ID	Platform	Capability
1	Accounting(0030c1-7fcc40)	HP J4812A ProCurve Switch...	S
2	Research(0060b0-889e43)	HP J4121A ProCurve Switch...	S
4	Support(0060b0-761a45)	HP J4121A ProCurve Switch...	S
7	Marketing(0030c5-38dc59)	HP J4813A ProCurve Switch...	S
12	Mgmt NIC(099a05-09df9b)	NIC Model X666	H
12	Mgmt NIC(099a05-09df11)	NIC Model X666	H

Figure 87. Example of CDP Neighbors Table Listing

Figure 88 illustrates a topology of CDP-enabled devices for the CDP Neighbors table listing in figure 87.

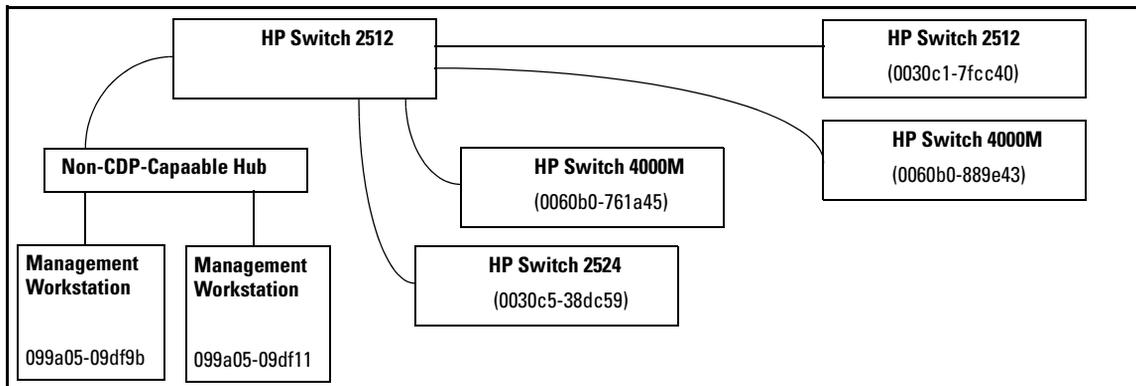


Figure 88. Example of CDP-Enabled Devices in a Topology for the Listing in Figure 87

Clearing (Resetting) the CDP Neighbors Table

This command removes any records of CDP neighbor devices from the switch’s CDP MIB objects.

Syntax: `cdp clear`

If you execute `cdp clear` and then execute `show cdp neighbors` before the switch receives a CDP packet from any neighbor device, the displayed table appears empty.

```

HP2512(config)# cdp clear
HP2512(config)# show cdp neighbors

CDP neighbors information

Port Device ID | Platform | Capability
-----+-----

```

Note that the table will again list entries after the switch receives new CDP packets from neighboring CDP devices.

Figure 89. View of the CDP Neighbors Table Immediately After Executing `cdp clear`

Configuring CDP Operation

Enabling or Disabling CDP Operation on the Switch. Enabling CDP operation (the default) on the switch causes the switch to:

- Transmit CDP packets describing itself to other, neighboring CDP devices
- 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, prevents the switch from transmitting outbound CDP packets to advertise itself to neighboring CDP devices, 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

For example, to disable CDP on the switch:

```
HP2512(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 : No
```

Enabling or Disabling CDP Operation on Individual Ports. In the factory-default configuration, the switch has all ports enabled and transmitting CDP packets. Disabling CDP on a port prevents that port from sending outbound CDP packets and causes it to drop inbound CDP packets without recording their data in the CDP Neighbors table. Suppose, for example, that switches "A" and "B" in figure 90 are running CDP, and that port 1 on switch "A" is connected to port 5 on switch "B". If you disable CDP on port 1 of switch "A", then switch "B" will no longer receive CDP packets from switch "A" and switch "A" will drop the CDP packets it receives from switch "B".

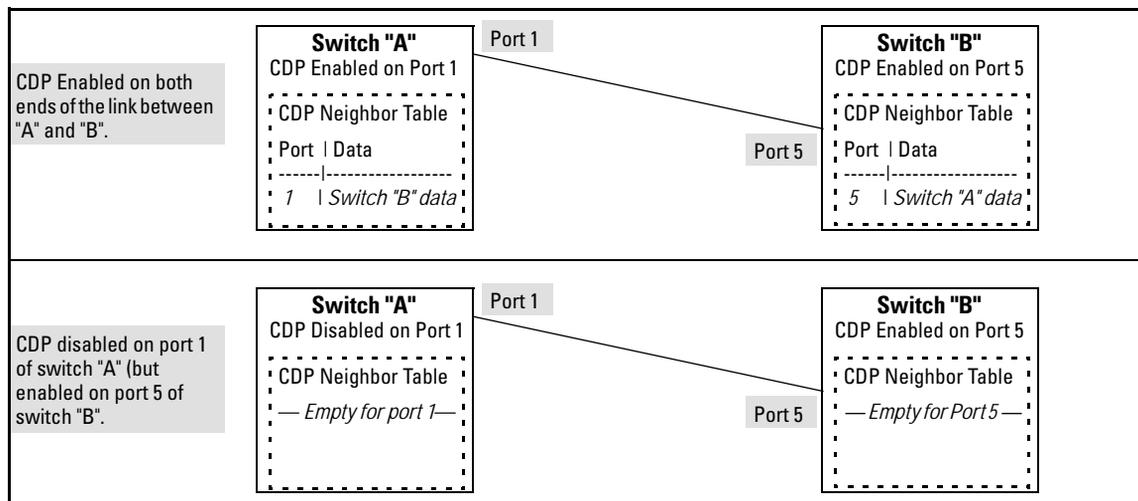


Figure 90. Example of Disabling CDP on an Individual Port

(The switch "A" entry in the switch "B" CDP Neighbors table remains until the **cdp holdtime** (time-to-live; set in switch "A") expires. Until then, the **show cdp neighbors** command continues to list switch "A" on port 5 of switch "B".)

Syntax: [no] cdp enable < [ethernet] port-list >

For example, to disable CDP on port 1 of a switch:

```
HP2512(config) no cdp enable 1
```

Changing the Transmission Interval for Outbound CDP Packets. The default interval the switch uses to transmit CDP packets describing itself to other, neighbor devices is 60 seconds. This command changes the interval.

Syntax: cdp timer < 5.. 254 >

For example, to reset a switch's transmit interval for CDP packets to one minute:

```
HP2512(config) cdp timer 60
```

Changing the Hold Time (CDP Packet Time-To-Live) for a Switch's CDP Packet

Information. The default hold time for the switch's CDP packet information in the CDP Neighbors table of another CDP device is 180 seconds (range: 5 - 254). This parameter is controlled in the transmitting switch, and applies to all outbound CDP packets the switch transmits.

Syntax: `cdp holdtime < 5 . . 254 >`

For example, to configure a switch's outbound CDP packets to live for one minute in the CDP Neighbors table of neighboring CDP devices:

```
HP2512(config) cdp holdtime 60
```

Effect of Spanning Tree (STP) On CDP Packet Transmission

If STP has blocked a port on the switch, that port does not transmit CDP packets. However, the port still receives CDP packets if the device on the other end of the link has CDP enabled. Thus, for example, if switch "A" has two ports linked to switch "B" (a CDP neighbor and the STP root device) and STP blocks traffic on one port and forwards traffic on the other:

- Switch "A" sends outbound CDP packets on the forwarding link, and the switch "B" CDP Neighbors table shows switch "A" on only one port.
- Switch "B" sends outbound CDP packets on both links, and the switch "A" CDP Neighbors table shows switch "B" on both ports.

To summarize, in a CDP neighbor pair running STP with redundant links, if one of the switches is the STP root, it transmits CDP packets out all ports connecting the two switches, while the other switch transmits CDP packets out only the unblocked port. Thus, the STP root switch will appear on multiple ports in the non-root switches CDP Neighbors table, while the non-root switch will appear on only one port in the root switch's CDP Neighbors table.

How CDP Selects the CDP Neighbor's IP Address When Multiple VLANs Are Present

When a switch detects a CDP neighbor and there are multiple VLANs configured on the neighbor's port, the switch uses the following criteria to determine which IP address to use when listing the neighbor in the CDP Neighbor table:

1. If only one VLAN on the neighbor's port has an IP address, the switch uses that IP address.
2. If the Primary VLAN on the neighbor's port has an IP address, the switch uses the neighbor's Primary VLAN IP address.

Enhancements in Release F.02.02

3. If 1 and 2 do not apply, then the switch determines which VLANs on the neighbor's port have IP addresses and uses the IP address of the VLAN with the lowest VID (VLAN Identification number) in this group.
4. If a CDP switch does not detect an IP address on the connecting port of a CDP neighbor, then the loopback IP address is used (127.0.0.1).

For example, in figure 91, port 1 on CDP switch "X" is connected to port 5 on CDP neighbor switch "Y", with the indicated VLAN configuration on port 5:

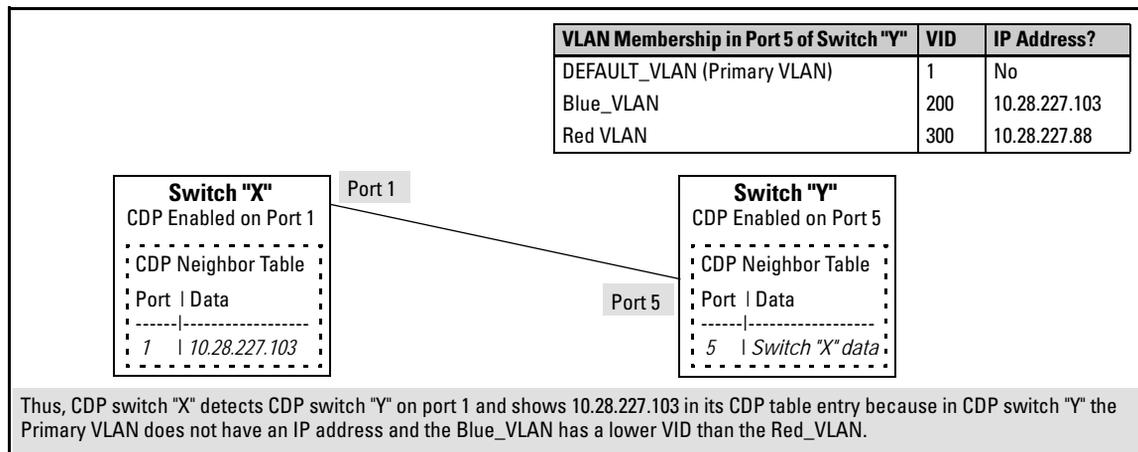


Figure 91. Example of IP Address Selection when the CDP Neighbor Has Multiple VLANs with IP Addresses

CDP Neighbor Data and MIB Objects

The switch places the data received from inbound CDP packets into its MIB (Management Information Base). This data is available in three ways:

- Using the switch's **show cdp neighbors** command to display a subset of Neighbor data
- Using the **walkmib** command to display a listing of the CDP MIB objects
- Electronically, using an SNMP utility designed to search the MIB for CDP data

As shown under "Viewing the Current Contents of the Switch's CDP Neighbors Table" on page 183, you can list a subset of data for each CDP device currently found in the switch's CDP Neighbors table. Table 16, "CDP Neighbors Data", describes the CDP Neighbor data set available in the Series 2500 switches.

Table 16. CDP Neighbors Data

CDP Neighbor Data	Displayed Neighbors Table	MIB	
Address Type	No	Yes	Always "1" (IP address only).
CDP Cache Address	No	Yes	IP address of source device.
Software Version	Yes	Yes	ASCII String
Device Name (ASCII string)	Yes	Yes	In HP ProCurve switches, this is the value configured for the System Name parameter.
Device MAC Address	Yes	Yes	Included in the Device Name entry.
Destination Port Number	Yes	Yes	On the Series 2500 switch (the receiving device), the number of the port through which the CDP packet arrived.
Source Port Number	No	Yes	On the source (neighbor) device, the number of the port through which the CDP packet was sent.
Product Name (ASCII string)	Yes	Yes	Platform name designated by vendor.
Capability Code (Device Type)	Yes (alpha character)	Yes (numeric character)	1 or R: Router 2: Transparent Bridge 4 or B: Source Route Bridge 8 or S: Switch 16 or H: Host 32 or I: IGMP conditional filtering 64 or r: Repeater

Displaying CDP Neighbor Data. To display the superset of CDP neighbor data held in the MIB, use the **walkmib** command.

Syntax: walkmib < MIB-identifier >

For example, with only one CDP device connected to the switch, you would see a **walkmib** listing similar to this:

```

HP2512(config)# walkmib CdpCacheEntry
cdpCacheAddressType.1.3 = 1
cdpCacheAddressType.2.3 = 1
cdpCacheAddress.1.3 = 0a 1c e3 66
cdpCacheVersion.1.3 = Revision C.09.X1 /sw/code/build/vgro(v00)
cdpCacheDeviceId.1.3 = North Campus 1(000080-000000)
cdpCacheDevicePort.1.3 = A1
cdpCachePlatform.1.3 = HP J4121A ProCurve Switch 4000M
cdpCacheCapabilities.1.3 = 8

```

Figure 92. Example of CDP Neighbor Data in the Series 2500 Switch MIB

For the current Series 2500 switch MIB, go to the **technical support** area at <http://www.hp.com/go/hpprocurve>.

CDP Operating Notes

Neighbor Maximum. The Series 2500 switches support up to 60 neighbors in the CDP Neighbors table. Even though the switches offer only 12 or 24 ports, multiple CDP devices can be neighbors on the same port if they are connected to the switch through a non-CDP device.

CDP Version Data. The Series 2500 switches use CDP-V1, but do not include IP prefix information, which is a router function; not a switch application.

Port Trunking with CDP. Where a static or LACP trunk forms the link between the switch and another CDP device, only one physical link in the trunk is used to transmit outbound CDP packets.

CDP-Capable Hubs. Some hubs are capable of running CDP, but also forward CDP packets as if the hub itself were transparent to CDP. Such hubs will appear in the switch's CDP Neighbor table and will also maintain a CDP neighbor table similar to that for switches. For more information, refer to the documentation provided for the specific hub.

Troubleshooting CDP Operation

The switch does not appear in the CDP Neighbors table of an adjacent CDP Device. This may be due to any of the following:

- Either the port connecting the switch to the adjacent device is not a member of an untagged VLAN or any Untagged VLAN to which the port belongs does not have an IP address.
- If there is more than one physical path between the switch and the other CDP device and STP is running on the switch, then STP will block the redundant link(s). In this case, the switch port on the remaining open link may not be a member of an untagged VLAN, or any untagged VLANs to which the port belongs may not have an IP address.
- The adjacent device's CDP Neighbors table may be full. Refer to the documentation provided for the adjacent CDP device to determine the table's capacity, and then view the device's Neighbors table to determine whether it is full.

One or more CDP neighbors appear intermittently or not at all in the switch's CDP Neighbors table. This may be caused by more than 60 neighboring devices sending CDP packets to the switch. Exceeding the 60-neighbor limit can occur, for example, where multiple neighbors are connected to the switch through non-CDP devices such as many hubs.

The Same CDP Switch or Router Appears on More Than One Port in the CDP Neighbors Table. Where CDP is running, a switch or router that is the STP root transmits outbound CDP packets over all links, including redundant links that STP may be blocking in non-root devices. In this case, the non-root device shows an entry in its CDP Neighbors table for every port on which it receives a CDP packet from the root device. See “Effect of Spanning Tree (STP) On CDP Packet Transmission” on page 187.

New Time Synchronization Protocol Options

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.

Formerly, TimeP was the only time protocol available for time synchronization in Series 2500 switches. Beginning with software release F.02.02, the switches also offer SNTP (Simple Network Time Protocol) and a new **timesync** command for changing the time protocol selection (or turning off time protocol operation).

Notes

- Although you can create and save configurations for both time protocols without conflicts, the switch allows only one active time protocol at any time.
- Time synchronization is no longer active in the factory default configuration. You must first select the desired protocol (default: TimeP), and then enable it.
- In the factory-default configuration for release F.02.02 and later, the time synchronization method is set to TimeP, with actual TimeP operation disabled. (In earlier releases, TimeP was enabled with DHCP for acquiring a TimeP server address).
- If you configure SNTP operation in the switch, but later download a configuration created using a pre-F.02.02 version of the software, the SNTP configuration will be replaced by the non-SNTP time synchronization settings in the downloaded configuration file.
- In the menu interface, the time protocol parameters have been moved from the "Internet (IP) Service" screen to the "System Information" screen, and the menu path is now:

2. Switch Configuration...

1. System Information

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. Refer to 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.
-

Overview: Selecting a Time Synchronization Protocol or Turning Off Time Protocol Operation

General Steps for Running a Time Protocol on the Switch:

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.

Note that 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:

- In the System Information screen of the Menu interface, set the **Time Synch Method** parameter to **None**, then press **[Enter]**, then **[S]** (for **Save**).
- In the config level of the CLI, execute **no timesync**.

SNTP: Viewing, Selecting, and Configuring

SNTP Features

Feature	Default	Menu	CLI	Web
view the SNTP time synchronization configuration	n/a	page 195	page 198	—
select SNTP as the time synchronization method	timep	page 196	pages 198 ff.	—
disable time synchronization	timep	page 196	page 201	—
enable the SNTP mode (Broadcast, Unicast, or Disabled)	disabled			—
broadcast	n/a	page 196	page 199	—
unicast	n/a	page 196	page 199	—
none/disabled	n/a	page 196	page 202	—
configure an SNTP server address (for Unicast mode only)	none	page 196	pages 199 ff.	—
change the SNTP server version (for Unicast mode only)	3	page 197	page 201	—
change the SNTP poll interval	720 seconds	page 197	page 201	—

Table 17. 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 timesync 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, it the the switch accepts a broadcast time update from the next server it detects.
Poll Interval (seconds)	In Unicast Mode: Specifies how often the switch polls the designated SNTP server for a time update. In Broadcast Mode: Specifies how often the switch polls the network broadcast address for a time update.
Server Address	Used only when the SNTP Mode is set to Unicast . Specifies the IP address of the SNTP server that the switch accesses for time synchronization updates. You can configure up to three servers; one using the menu or CLI, and two more using the CLI. See “SNTP Unicast Time Polling with Multiple SNTP Servers” on page 209.
Server Version	Default: 3; range: 1 - 7. 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.

Menu: Viewing and Configuring SNTP

To View, Enable, and Modify SNTP Time Protocol:

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

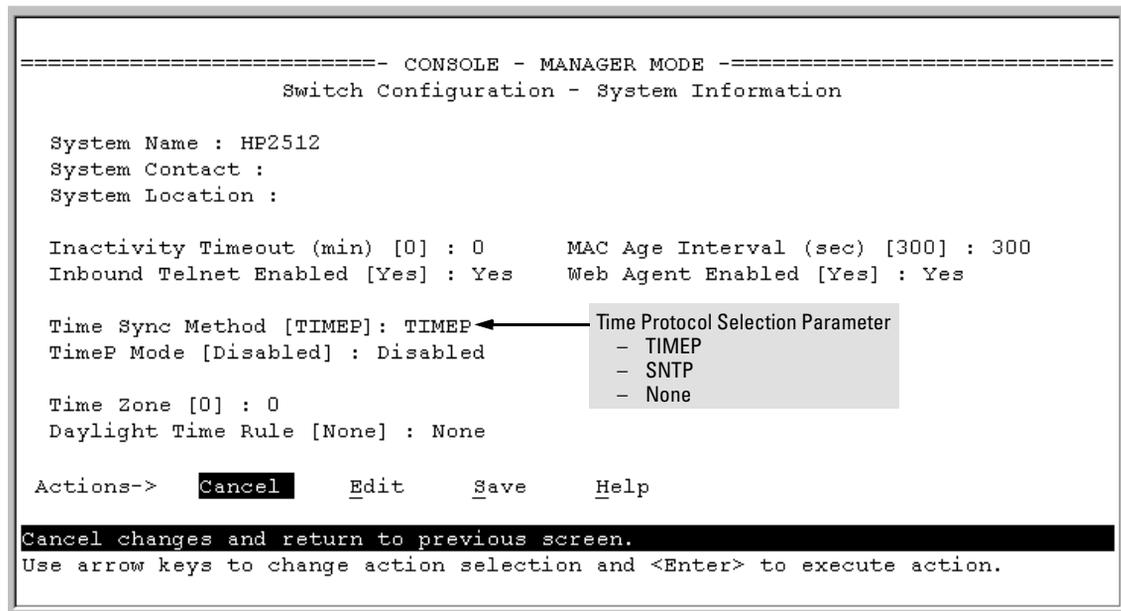
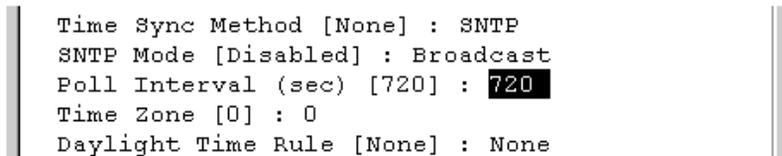


Figure 93. The System Information Screen (Default Values)

2. Press [E] (for **Edit**). The cursor moves to the **System Name** field.
3. Use to move the cursor to the **Time Sync Method** field.
4. Use the Space bar to select **SNTP**, then press once to display and move to the **SNTP Mode** field.
5. Do one of the following:
 - Use the Space bar to select the **Broadcast** mode, then press to move the cursor to the **Poll Interval** field, and go to step 6. (For Broadcast mode details, see "SNTP Operating Modes" on page 193.)



- Use the Space bar to select the **Unicast** mode, then do the following:
 - i. Press to move the cursor to the **Server Address** field.
 - ii. 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), then see “SNTP Unicast Time Polling with Multiple SNTP Servers” on page 209.

- iii. Press to move the cursor to the **Server Version** field. Enter the value that matches the SNTP server version running on the device you specified in the preceding step (step ii). 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 causes the switch to delete the primary SNTP server from the server list and to select a new primary SNTP server from the IP address(es) in the updated list. For more on this topic, see “SNTP Unicast Time Polling with Multiple SNTP Servers” on page 209.

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

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

6. In the **Poll Interval** field, enter the time in seconds that you want for a Poll Interval. (For Poll Interval operation, see table 17, “SNTP Parameters”, on page 195.)
7. Press **[Enter]** to return to the Actions line, then **[S]** (for **S**ave) to enter the new time protocol configuration in both the startup-config and running-config files.

CLI: Viewing and Configuring SNTP

CLI Commands Described in this Section

show sntp	page 198
[no] timesync	pages 198 and ff., 201
sntp broadcast	page 199
sntp unicast	page 199
sntp server	pages 199 and ff.
Protocol Version	page 201
poll-interval	page 201
no sntp	page 202

This section describes how to use the CLI to view, enable, and configure SNTP parameters.

Viewing the Current SNTP Configuration

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

Syntax: show sntp

For example, if you configured the switch with SNTP as the time synchronization method, then enabled SNTP in broadcast mode with the default poll interval, **show sntp** lists the following:

```
HP2512# show sntp
SNTP Configuration
Time Sync Mode: Sntp
SNTP Mode : Broadcast
Poll Interval (sec) [720] : 720
```

Figure 94. Example of SNTP Configuration When SNTP Is the Selected Time Synchronization Method

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. For example:

```
HP2512# show sntp
SNTP Configuration
Time Sync Mode: Timep
SNTP Mode : Broadcast
Poll Interval (sec) [720] : 720
```

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

Figure 95. Example of SNTP Configuration When SNTP Is Not the Selected Time Synchronization Method

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
sntp < broadcast | unicast >
sntp server < ip-addr >
sntp poll-interval < 30 .. 720 >

*Selects SNTP as the time protocol.
Enables the SNTP mode (below and page 199).
Required only for unicast mode (page 199).
Enabling the SNTP mode also enables the SNTP poll interval (default: 720 seconds; page 201).*

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.*
 sntp broadcast *Configures **Broadcast** as the SNTP mode.*

For example, suppose:

- Time synchronization is in the factory-default configuration (TimeP is the currently selected time synchronization method).
- You want to:
 1. View the current time synchronization.
 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:

```
HP2512(config)# show sntp ❶
SNTP Configuration
  Time Sync Mode: Timep
  SNTP Mode : disabled
  Poll Interval (sec) [720] : 720
HP2512(config)# timesync sntp ❷
HP2512(config)# sntp broadcast ❸
HP2512(config)# show sntp ❹
SNTP Configuration
  Time Sync Mode: Sntp
  SNTP Mode : Broadcast
  Poll Interval (sec) [720] : 720
```

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

❹ *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.

Figure 96. Example of Enabling SNTP Operation in Broadcast Mode

Enabling SNTP in Unicast Mode. Like broadcast mode, configuring SNTP for unicast mode enables SNTP. However, for Unicast operation, you must also specify the IP address of at least one SNTP server. The switch allows up to three unicast servers. You can use the Menu interface or the CLI to configure one server or to replace an existing Unicast server with another. To add a second or third server, you must use the CLI. For more on SNTP operation with multiple servers, see “SNTP Unicast Time Polling with Multiple SNTP Servers” on page 209.

Syntax: timesync sntp *Selects SNTP as the time synchronization method.*
snmp unicast *Configures the SNTP mode for Unicast operation.*
snmp server <ip-addr> [version] *Specifies the SNTP server. The default server version is 3.*
no snmp server <ip-addr> *Deletes the specified SNTP server.*

Note

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

For 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):

```
HP2512(config)# timesync sntp Selects SNTP.  
HP2512(config)# snmp unicast Activates SNTP in Unicast mode.  
HP2512(config)# snmp server 10.28.227.141 Specifies the SNTP server and  
accepts the current SNTP server  
version (default: 3)
```

```
HP2512(config)# show snmp  
SNTP Configuration  
Time Sync Mode: Sntp  
SNTP Mode : Unicast  
Poll Interval (sec) [720] : 720  
IP Address          Protocol Version  
-----  
10.28.227.141      3
```

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

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

Figure 97. Example of Configuring SNTP for Unicast Operation

If the SNTP server you specify uses SNTP version 4 or later, use the snmp server command to specify the correct version number. For example, suppose you learned that SNTP version 4 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 and then re-enter it with the correct version number for that server:

```

HP2512(config)# no sntp server 10.28.227.141
HP2512(config)# sntp server 10.28.227.141 4
HP2512(config)# show sntp
  SNTP Configuration
    Time Sync Mode: Sntp
    SNTP Mode : Broadcast
    Poll Interval (sec) [720] : 600

    IP Address          Protocol Version
    -----
    10.28.227.141      4

```

Deletes unicast SNTP server entry.

Re-enters the unicast server with a non-default protocol version.

show sntp displays the result.

Figure 98. Example of Specifying the SNTP Protocol Version Number

Changing the SNTP Poll Interval. This command lets you specify how long the switch waits between time polling intervals. The default is 720 seconds and the range is 30 to 720 seconds. (This parameter is separate from the poll interval parameter used for Timepeg operation.)

Syntax: sntp poll-interval < 30 .. 720 >

For example, to change the poll interval to 300 seconds:

```
HP2512(config)# sntp poll-interval 300
```

Disabling Time Synchronization Without Changing the SNTP Configuration. The recommended method for disabling time synchronization is to use the **timesync** command. This halts time synchronization without changing your SNTP configuration.

Syntax: no timesync

For 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:

```
HP2512(config)# no timesync
```

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

```

HP2524(config)# show sntp
  SNTP Configuration
    Time Sync Mode: Disabled
    SNTP Mode : Broadcast
    Poll Interval (sec) [720] : 720

```

Figure 99. Example of SNTP with Time Synchronization Disabled

Disabling the SNTP Mode. If you want to prevent SNTP from being used even if 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**.*

For 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.

```

HP2512(config)# no sntp
HP2512(config)# show sntp
SNTP Configuration
  Time Sync Mode: Sntp
  SNTP Mode : disabled
  Poll Interval (sec) [720] : 720
  IP Address          Protocol Version
  -----
  10.28.227.141      3
    
```

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

Figure 100. Example of Disabling Time Synchronization by Disabling the SNTP Mode

TimeP: Viewing, Selecting, and Configuring

Timep Features

Feature	Default	Menu	CLI	Web
view the Timep time synchronization configuration	n/a	page 203	page 205	—
select Timep as the time synchronization method	TIMEP	page 202	pages 206 ff.	—
disable time synchronization	timep	page 204	page 208	—
enable the Timep mode	Disabled			—
DHCP	—	page 204	page 206	—
manual	—	page 204	page 207	—
none/disabled	—	page 204	page 209	—
change the SNTP poll interval	720 minutes	page 205	page 208	—

Table 18. Timep Parameters

SNTP Parameter	Operation
Time Sync Method	Used to select either TIMEP (the default), SNTP, or None as the time synchronization method.
Timep Mode	
Disabled	The Default. Timep does not operate, even if specified by the Menu interface Time Sync Method parameter or the CLI timesync 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.
Poll Interval (minutes)	Default: 720 minutes. Specifies the interval the switch waits between attempts to poll the TimeP server for updates.

Menu: Viewing and Configuring TimeP

To View, Enable, and Modify the TimeP Protocol:

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

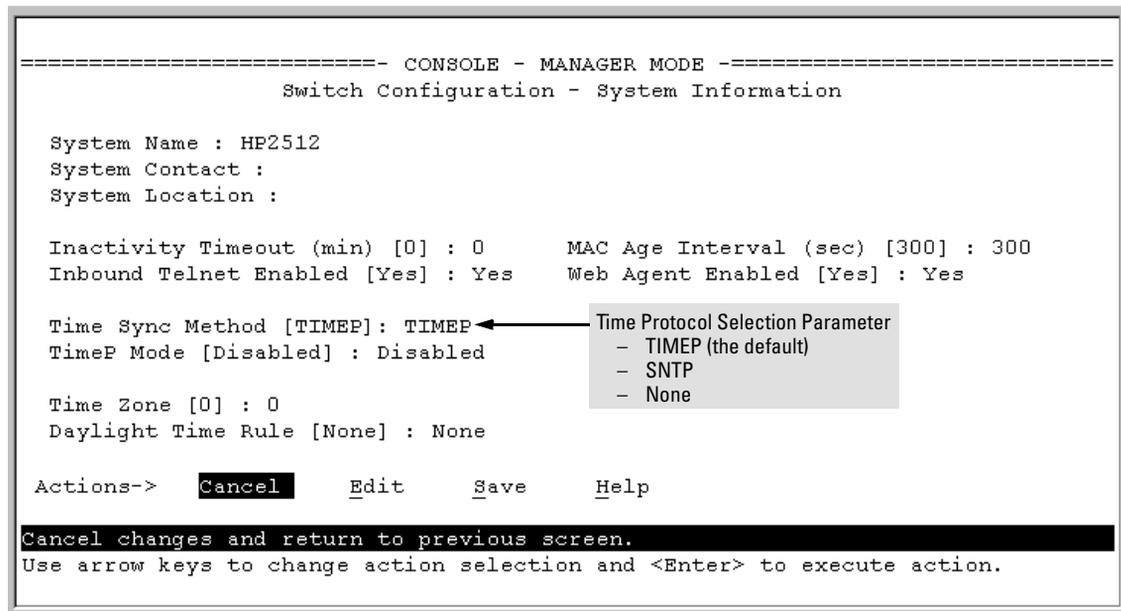
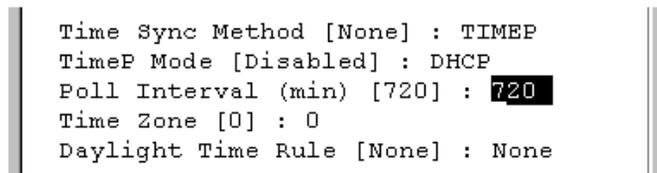


Figure 101. The System Information Screen (Default Values)

2. Press [E] (for **E**dit). The cursor moves to the **System Name** field.
3. Use to move the cursor to the **Time Sync Method** field.
4. If **TIMEP** is not already selected, use the Space bar to select **TIMEP**, then press once to display and move to the **TimeP Mode** field.
5. Do one of the following:
 - Use the Space bar to select the **DHCP** mode, then press to move the cursor to the **Poll Interval** field, and go to step 6.



- Use the Space bar to select the **Manual** mode.
 - i. Press to move the cursor to the **Server Address** field.
 - ii. 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.

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

```
Time Sync Method [None] : TIMEP
TimeP Mode [Disabled] : Manual      Server Address : 10.28.227.141
Poll Interval (min) [720] : 720
Time Zone [0] : 0
Daylight Time Rule [None] : None
```

6. In the **Poll Interval** field, enter the time in minutes that you want for a TimeP Poll Interval.

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.

CLI: Viewing and Configuring TimeP

CLI Commands Described in this Section

show timep	page 205
[no] timesync	page 206 ff., 208
ip timep	
dhcp	page 206
manual	page 207
server <ip-addr>	page 207
interval	page 208
no ip timep	page 209

This section describes how to use the CLI to view, enable, and configure TimeP parameters.

Viewing the Current TimeP Configuration

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

Syntax: show timep

For example, 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:

```
HP2512(config)# show timep
Timep Configuration
Time Sync Mode: Timep
TimeP Mode : DHCP      Poll Interval (min) : 720
```

Figure 102. Example of TimeP Configuration When TimeP Is the Selected Time Synchronization Method

If SNTP is the selected time synchronization method), **show timep** still lists the TimeP configuration even though it is not currently in use:

```
HP2512(config)# show timep
Timep Configuration
Time Sync Mode: Sntp
TimeP Mode : DHCP Poll Interval (min) : 720
```

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

Figure 103. Example of SNTP Configuration When SNTP Is Not the Selected Time Synchronization Method

Configuring (Enabling or Disabling) the TimeP Mode

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

Syntax:	<code>timesync timep</code>	<i>Selects TimeP as the time protocol.</i>
	<code>ip timep < dhcp manual ></code>	<i>Enables the selected TimeP mode.</i>
	<code>no ip timep</code>	<i>Disables the TimeP mode.</i>
	<code>no timesync</code>	<i>Disables the time protocol.</i>

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:	<code>timesync timep</code>	<i>Selects TimeP as the time synchronization method.</i>
	<code>ip timep dhcp</code>	<i>Configures DHCP as the TimeP mode.</i>

For example, suppose:

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

The commands and output would appear as follows:

```

HP2512 (config)# show timep 1 show timep displays the TimeP configuration and also shows
Timep Configuration that SNTP is the currently active time synchronization mode.
Time Sync Mode: Sntp
TimeP Mode : Disabled

HP2512 (config)# timesync timep 2

HP2512 (config)# ip timep dhcp 3

HP2512 (config)# show timep 4 show timep again displays the TimeP configuration and shows that TimeP is
Timep Configuration now the currently active time synchronization mode.
Time Sync Mode: Timep
TimeP Mode : DHCP   Poll Interval (min) : 720

```

Figure 104. Example of Enabling TimeP Operation in DHCP Mode

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	<i>Selects Timep.</i>
	ip timep manual <ip-addr>	<i>Activates TimeP in Manual mode with a specified TimeP server.</i>
	no ip timep	<i>Disables TimeP.</i>

Note

To change from one TimeP server to another, you must (1) use the **no ip timep** command to disable TimeP mode, and then reconfigure TimeP in Manual mode with the new server IP address.

For 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):

```
HP2512(config)# timesync timep Selects TimeP.  
HP2512(config)# ip timep manual 10.28.227.141 Activates TimeP in Manual mode.
```

```
HP2512(config)# timesync timep  
HP2512(config)# ip timep manual 10.28.227.141  
  
HP2512(config)# Show timep  
Timep Configuration  
Time Sync Mode: Timep  
TimeP Mode : Manual          Server Address : 10.28.227.141  
Poll Interval (min) : 720
```

Figure 105. Example of Configuring Timep for Manual Operation

Changing the TimeP Poll Interval. This command lets you specify 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.)

Syntax: ip timep dhcp interval < 1 .. 9999 >
ip timep manual interval < 1 .. 9999 >

For example, to change the poll interval to 60 minutes:

```
HP2512(config)# ip timep interval 60
```

Disabling Time Synchronization Without Changing the TimeP Configuration. The recommended method for disabling time synchronization is to use the **timesync** command. This halts time synchronization without changing your TimeP configuration.

Syntax: no timesync

For 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:

```
HP2512(config)# no timesync
```

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

```
HP2524(config)# show timep  
Timep Configuration  
Time Sync Mode: Disabled  
TimeP Mode : DHCP          Poll Interval (min) : 720
```

Figure 106. Example of TimeP with Time Synchronization Disabled

Disabling the TimeP Mode. Disabling the TimeP mode means to configure it as disabled. (Disabling TimeP prevents the switch from using it as the time synchronization protocol, even if it is the selected **Time Sync Method** option.)

Syntax: no ip timep *Disables TimeP by changing the TimeP mode configuration to Disabled.*

For example, if the switch is running TimeP in DHCP mode, **no ip timep** changes the TimeP configuration as shown below, and disables time synchronization on the switch.

```
HP2512(config)# no ip timep

HP2512(config)# show timep
Timep Configuration
  Time Sync Mode: Timep
  TimeP Mode : Disabled
```

Even though the Time Sync Mode is set to Timep, time synchronization is disabled because no ip timep has disabled the TimeP Mode parameter.

Figure 107. Example of Disabling Time Synchronization by Disabling the TimeP Mode Parameter

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.

Address Prioritization

If you use the CLI to configure multiple SNTP servers, the switch prioritizes them according to the decimal values of their IP addresses. That is, the switch compares the decimal value of the octets in the addresses and orders them accordingly, with the lowest decimal value assigned as the primary address, the second-lowest decimal value assigned as the next address, and the third-lowest decimal value as the last address. If the first octet is the same between two of the addresses, the second octet is compared, and so on. For example:

SNTP Server IP Address	Server Ranking According to Decimal Value of IP Address
10.28.227.141	Primary
10.28.227.153	Secondary
10.29.227.100	Tertiary

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. For example, suppose you have already configured the primary address in the above table (10.28.227.141). To configure the remaining two addresses, you would do the following:

```
HP2512(config)# sntp server 10.29.227.100
HP2512(config)# sntp server 10.28.227.153
HP2512(config)# show sntp
SNTP Configuration
Time Sync Mode: Sntp
SNTP Mode : disabled
Poll Interval (sec) [720] : 720
IP Address          Protocol Version
-----
10.28.227.141      3
10.28.227.153      3
10.29.227.100      3
```

Prioritized list of SNTP Server IP Addresses

Figure 108. Example of SNTP Server Address Prioritization

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. To delete an address, you must use the CLI. If there are multiple addresses and you delete one of them, the switch re-orders the address priority. (See “Address Prioritization” on page 209.)

Syntax: no sntp server < ip-addr >

For example, to delete the primary address in the above example (and automatically convert the secondary address to primary):

```
HP2512(config)# no sntp server 10.28.227.141
```

Menu Interface Operation with Multiple SNTP Server Addresses Configured

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. If there are multiple addresses configured, the switch re-orders the addresses according to the criteria described under “Address Prioritization” on page 209. For example, suppose the switch already has the following three SNTP server IP addresses configured.

- 10.28.227.141 (primary)
- 10.28.227.153 (secondary)
- 10.29.227.100 (tertiary)

If you use the Menu interface to add 10.28.227.160, the new prioritized list will be:

New Address List	Address Status
10.28.227.153	New Primary (The former primary, 10.28.227.141 was deleted when you used the menu to add 10.28.227.160.)
10.28.227.160	New Secondary
10.29.227.100	Same Tertiary (This address still has the highest decimal value.)

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.

Operation and Enhancements for Multimedia Traffic Control (IGMP)

How Data-Driven IGMP Operates

The information in this section supplements the information provided under "Multimedia Traffic Control with IP Multicast (IGMP)" beginning on page 9-91 in the Management and Configuration Guide included with your Series 2500 switch and also available at <http://www.hp.com/go/hpprocurve>.

This section uses the following terms to describe IGMP operation:

- **Querier:** A required IGMP device that facilitates the IGMP protocol and traffic flow on a given LAN. This device tracks which ports are connected to devices (IGMP clients) that belong to specific multicast groups, and triggers updates of this information. With IGMP enabled, the Series 2500 switches use data from the Querier to determine whether to forward or block multicast traffic on specific ports. When the switch has an IP address on a given VLAN, it automatically operates as a Querier for that VLAN if it does not detect a multicast router or another switch functioning as a Querier.
- **IGMP Device:** A switch or router running IGMP traffic control features.
- **IGMP Host:** An end-node device running an IGMP (multipoint, or multicast communication) application.

Without IGMP enabled, the switch simply floods all IP multicast traffic it receives on a given VLAN through all ports on that VLAN (except the port on which it received the traffic). This can result in significant and unnecessary bandwidth usage in networks where IP multicast traffic is a factor. Data-Driven IGMP reduces this problem by authorizing the switch to restrict multicast traffic only to ports where a given multicast group should flow.

Series 2500 switches (all software versions) use data-driven IGMP to better control IP multicast traffic.

An IP multicast packet includes the multicast group (address) to which the packet belongs. When an IGMP client connected to a switch port needs to receive multicast traffic from a specific group, it joins the group by sending an IGMP report (join request) to the network. (The multicast group specified in the join request is determined by the requesting application running on the IGMP client.) When a networking device with IGMP enabled receives the join request for a specific group, it forwards any IP multicast traffic it receives for that group through the port on which the join request was received. To reduce unnecessary traffic, the networking device does not forward a given group's

multicast packets to ports from which a join request for that group has not been received. (If the switch or router has not received any join requests for a given multicast group, it drops the traffic it receives for that group.)

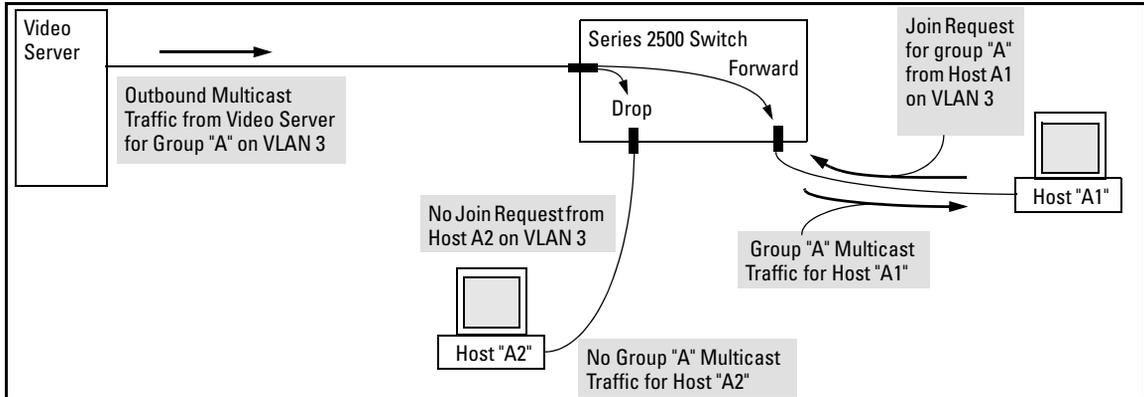


Figure 109. Example of Data-Driven IGMP Operation

Thus, after you enable IGMP on a VLAN configured in the switch, it continually listens for IGMP messages and IP multicast traffic on all ports in the VLAN, and forwards IGMP traffic for a given multicast address only through the port(s) on that VLAN where an IGMP report (join request) for that address was received from an IGMP client device.

Note

IP multicast traffic groups are identified by IP addresses in the range of 224.0.0.0 to 239.255.255.255.

Incoming IGMP packets intended for reserved, or "well-known" multicast addresses automatically flood through all ports (except the port on which the packets entered the switch). For more on this topic, see "The Switch Excludes Well-Known or Reserved Multicast Addresses from IP Multicast Filtering" on page 219.

IGMP Operates With or Without IP Addressing

Formerly, IGMP operation on the Series 2500 switches required an IP address and subnet mask for each VLAN running IGMP. Beginning with release F.02.02, you can configure IGMP on VLANs that do not have IP addressing. The benefit of IGMP without IP addressing is a reduction in the number of IP addresses you have to use and configure. This can be significant in a network with a large number of VLANs. The limitation on IGMP without IP addressing is that the switch cannot become Querier on any VLANs for which it has no IP address—so the network administrator must ensure that another IGMP device will act as Querier and that an additional IGMP device is available as a backup Querier.

IGMP Function Available With IP Addressing Configured on the VLAN	Available Without IP Addressing?	Operating Differences Without an IP Address
Drop multicast group traffic for which there have been no join requests from IGMP clients connected to ports on the VLAN.	Yes	None
Forward multicast group traffic to any port on the VLAN that has received a join request for that multicast group.	Yes	None
Forward join requests (reports) to the Querier.	Yes	None
Configure individual ports in the VLAN to Auto (the default)/ Blocked , or Forward .	Yes	None
Configure IGMP traffic forwarding to normal or high-priority forwarding.	Yes	None
Age-Out IGMP group addresses when the last IGMP client on a port in the VLAN leaves the group.	Yes	Requires that another IGMP device in the VLAN has an IP address and can operate as Querier. This can be a multi-cast router or another switch configured for IGMP operation. (HP recommends that the VLAN also include a device operating as a backup Querier in case the device operating as the primary Querier fails for any reason.
Support Fast-Leave IGMP (below) and Forced Fast-Leave IGMP (page 216).	Yes	
Support automatic Querier election.	No	Querier operation not available.
Operate as the Querier.	No	Querier operation not available.
Provide a backup Querier.	No	Querier operation not available.

Fast-Leave IGMP

IGMP Operation Presents a "Delayed Leave" Problem. Where multiple IGMP clients are connected to the same port on an IGMP device (switch or router), if only one IGMP client joins a given multicast group, then later sends a Leave Group message and ceases to belong to that group, the IGMP device retains that IGMP client in its IGMP table and continues forwarding IGMP traffic to the IGMP client until the Querier triggers confirmation that no other group members exist on the same port. This means that the switch continues to transmit unnecessary multicast traffic through the port until the Querier renews multicast group status.

Fast-Leave IGMP Reduces Leave Delays. Fast-Leave IGMP operates on a port if an IGMP client connects to the port and there are no other end nodes detected on that port. In this case, when the client leaves a multicast group, Fast-Leave IGMP automatically accelerates the blocking of further,

unnecessary multicast traffic from that group to the former IGMP client. This improves performance by reducing the amount of multicast traffic going through the port to the IGMP client after the client leaves a multicast group. IGMP in the Series 2500 switches automatically uses this Fast-Leave feature.

Automatic Fast-Leave Operation. If a Series 2500 switch port is :

- a. Connected to only one end node
- b. The end node currently belongs to a multicast group; i.e. is an IGMP client
- c. The end node subsequently leaves the multicast group

Then the switch does not need to wait for the Querier status update interval, but instead immediately removes the IGMP client from its IGMP table and ceases transmitting IGMP traffic to the client. (If the switch detects multiple end nodes on the port, automatic Fast-Leave does not activate—regardless of whether one or more of these end nodes are IGMP clients.)

In the next figure, automatic Fast-Leave operates on the switch ports for IGMP clients "3A" and "5B", but not on the switch port for IGMP clients "7A" and 7B, Server "7C", and printer "7D".

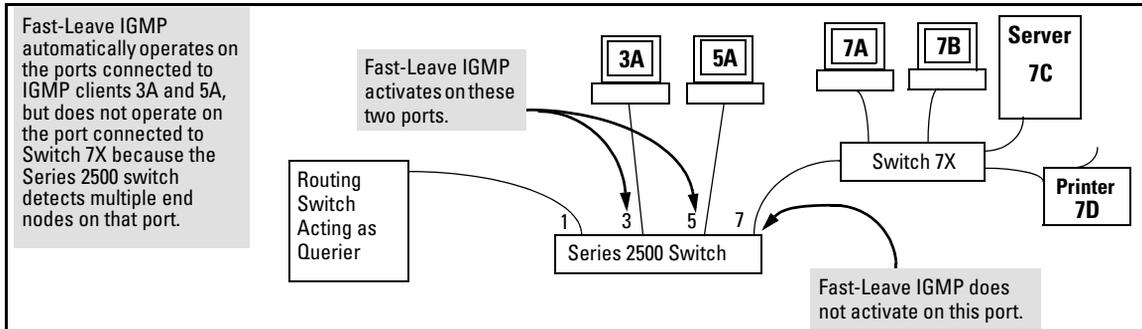


Figure 110. Example of Automatic Fast-Leave IGMP Criteria

When client "3A" running IGMP is ready to leave the multicast group, it transmits a Leave Group message. Because the Series 2500 switch knows that there is only one end node on port 3, it removes the client from its IGMP table and halts multicast traffic (for that group) to port 3. If the switch is not the Querier, it does not wait for the actual Querier to verify that there are no other group members on port 3. If the switch itself is the Querier, it does not query port 3 for the presence of other group members.

Note that Fast-Leave operation does not distinguish between end nodes on the same port that belong to different VLANs. Thus, for example, even if all of the devices on port 7 in figure 110 belong to different VLANs, Fast-Leave does not operate on port 7.

Forced Fast-Leave IGMP

Forced Fast-Leave IGMP Features

Feature	Default	Menu	CLI	Web
view the Forced Fast-Leave configuration				
view the switch's Forced Fast-Leave state	n/a	—	page 216	—
configure Forced Fast-Leave				
configure Forced Fast-Leave for an individual port	2 (disabled)	—	page 218	—

Forced Fast-Leave IGMP speeds up the process of blocking unnecessary IGMP traffic to a switch port that is connected to multiple end nodes. (This feature does not activate on ports where the switch detects only one end node. Instead, the regular Fast Leave described in the preceding section activates.) For example, in figure 110, even if you configured Forced Fast-Leave on all ports in the switch, the feature would activate only on port 7 (which has multiple end nodes) when a Leave Group request arrived on that port.

When a port having multiple end nodes receives a Leave Group request from one end node for a given multicast group "X", Forced Fast-Leave activates and waits a small amount of time to receive a join request from any other group "X" member on that port. If the port does not receive a join request for that group within the forced-leave interval, the switch then blocks any further group "X" traffic to the port.

Configuration Options for Forced Fast-Leave

Feature	Default	Settings	Function
Forced Fast-Leave state	2 (disabled)	1 (enabled) 2 (disabled)	Uses the setmib command to enable or disable Forced Fast-Leave on individual ports. When enabled on a port, Forced Fast-Leave operates only if the switch detects multiple end nodes (and at least one IGMP client) on that port.

CLI: Listing the Forced Fast-Leave Configuration

The Forced Fast-Leave configuration includes the state (enabled or disabled) for each port and the Forced-Leave Interval for all ports on the switch.

To list the Forced Fast-Leave state for all ports in the switch:

```
Syntax:    HP2512# walkmib hpSwitchIgmpportForcedLeaveState.1
              or
              HP2512# walkmib 1.3.6.1.4.1.11.2.14.11.5.1.7.1.15.3.1.5.1
```

For example:

```
HP2512# walkmib hpswitchigmpportforcedleavestate.1
hpSwitchIgmpPortForcedLeaveState.1.1 = 2
hpSwitchIgmpPortForcedLeaveState.1.2 = 2
hpSwitchIgmpPortForcedLeaveState.1.3 = 2
hpSwitchIgmpPortForcedLeaveState.1.4 = 2
hpSwitchIgmpPortForcedLeaveState.1.5 = 2
hpSwitchIgmpPortForcedLeaveState.1.6 = 2
hpSwitchIgmpPortForcedLeaveState.1.7 = 2
hpSwitchIgmpPortForcedLeaveState.1.8 = 2
hpSwitchIgmpPortForcedLeaveState.1.9 = 2
hpSwitchIgmpPortForcedLeaveState.1.10 = 2
hpSwitchIgmpPortForcedLeaveState.1.11 = 2
hpSwitchIgmpPortForcedLeaveState.1.12 = 2
hpSwitchIgmpPortForcedLeaveState.1.13 = 2
hpSwitchIgmpPortForcedLeaveState.1.14 = 2
```

In this example, the **2** at the end of each port listing shows that Fast Forced-Leave is disabled on all ports in the switch.

Figure 111. Listing the Forced Fast-Leave State for Ports in an HP2512 Switch

To list the Forced Fast-Leave state for a single port.

Syntax: getmib hpSwitchIgmpPortForcedLeaveState.1. <port-number> (Not case-sensitive.)
 getmib 1.3.6.1.4.1.11.2.14.11.5.1.7.1.15.3.1.5.1. <port-number>

For example, to use either of the above command versions to list the state for port 7:

```
HP2512# getmib hpswitchigmpportforcedleavestate.1.7
hpSwitchIgmpPortForcedLeaveState.1.7 = 2
```

The **7** specifies port 7.
 The **2** shows that Fast Forced-Leave is disabled on port 7.

Figure 112. Listing the Forced Fast-Leave State for a Single Port

CLI: Configuring Per-Port Forced Fast-Leave IGMP

In the factory-default configuration, Forced Fast-Leave is disabled for all ports on the switch. To enable (or disable) this feature on individual ports, use the switch's MIB commands, as shown below.

Syntax: setmib hpSwitchIgmpportForcedLeaveState.1.<port-number> -i < 1 | 2 >
 or
 setmib 1.3.6.1.4.1.11.2.14.11.5.1.7.1.15.3.1.5.1.<port-number> -i < 1 | 2 >

where:

1 = Forced Fast-Leave enabled

2 = Forced Fast-Leave disabled

For example, to enable Forced Fast-Leave on ports 7 and 8:

```

HP2512# setmib hpSwitchIgmpportForcedLeaveState.1.7 -i 1
hpSwitchIgmpportForcedLeaveState.1.7 = 1
HP2512# setmib hpSwitchIgmpportForcedLeaveState.1.8 -i 1
hpSwitchIgmpportForcedLeaveState.1.8 = 1
  
```

Figure 113. Example of Changing the Forced Fast-Leave Configuration on Ports 7 and 8

Querier Operation

The function of the IGMP Querier is to poll other IGMP-enabled devices in an IGMP-enabled VLAN to elicit group membership information. The switch performs this function if there is no other device in the VLAN, such as a multicast router, to act as Querier. Although the switch automatically ceases Querier operation in an IGMP-enabled VLAN if it detects another Querier on the VLAN, you can also use a CLI command to disable the Querier function for that VLAN. For example, to disable the Querier function on VLAN 1 in a Series 2500 switch:

```

HP2512(config)# no vlan 1 ip igmp querier     Disables Querier function on VLAN 1.
HP2512(vlan-1)# no ip igmp querier         Disables Querier function on VLAN 1
                                              from within the VLAN 1 context.
  
```

Note

A Querier is required for proper IGMP operation. For this reason, if you disable the Querier function on a Series 2500 switch, ensure that there is an IGMP Querier (and, preferably, a backup Querier) available on the same VLAN.

If the switch becomes the Querier for a particular VLAN (for example, the DEFAULT_VLAN), then subsequently detects queries transmitted from another device on the same VLAN, the switch ceases to operate as the Querier for that VLAN. If this occurs, the switch Event Log lists a pair of messages similar to these:

```
I 01/15/01 09:01:13 igmp: DEFAULT_VLAN: Other Querier detected
I 01/15/01 09:01:13 igmp: DEFAULT_VLAN: This switch is no longer Querier
```

In the above scenario, if the other device ceases to operate as a Querier on the default VLAN, then the switch detects this change and can become the Querier as long as it is not pre-empted by some other IGMP Querier on the VLAN. In this case, the switch Event Log lists messages similar to the following to indicate that the switch has become the Querier on the VLAN:

```
I 01/15/01 09:21:55 igmp: DEFAULT_VLAN: Querier Election in process
I 01/15/01 09:22:00 igmp: DEFAULT_VLAN: This switch has been elected as Querier
```

The Switch Excludes Well-Known or Reserved Multicast Addresses from IP Multicast Filtering

Each multicast host group is identified by a single IP address in the range of 224.0.0.0 through 239.255.255.255. Specific groups of consecutive addresses in this range are termed "well-known" addresses and are reserved for predefined host groups. IGMP does not filter these addresses, so any packets the switch receives for such addresses are flooded out all ports assigned to the VLAN on which they were received (except the port on which the packets entered the VLAN). The following table lists the 32 well-known address groups (8192 total addresses) that IGMP does not filter on the Series 2500 switches.

Table 19. Well-Known IP Multicast Address Groups Excluded from IGMP Filtering

Groups of Consecutive Addresses in the Range of 224.0.0.X to 239.0.0.X*		Groups of Consecutive Addresses in the Range of 224.128.0.X to 239.128.0.X*	
224.0.0.X	232.0.0.X	224.128.0.X	232.128.0.X
225.0.0.X	233.0.0.X	225.128.0.X	233.128.0.X
226.0.0.X	234.0.0.X	226.128.0.X	234.128.0.X
227.0.0.X	235.0.0.X	227.128.0.X	235.128.0.X
228.0.0.X	236.0.0.X	228.128.0.X	236.128.0.X
229.0.0.X	237.0.0.X	229.128.0.X	237.128.0.X
230.0.0.X	238.0.0.X	230.128.0.X	238.128.0.X

Groups of Consecutive Addresses in the Range of 224.0.0.X to 239.0.0.X*		Groups of Consecutive Addresses in the Range of 224.128.0.X to 239.128.0.X*	
231.0.0.x	239.0.0.x	231.128.0.x	239.128.0.x

* X is any value from 0 to 255.

Switch Memory Operation

If you are using the CLI to change the switch configuration, HP recommends that you use the **write memory** command to permanently save the changes (to the startup-config file) before exiting from the CLI. CLI configuration changes are not saved from the Menu interface to the startup-config file unless you make a configuration change in the Menu interface before using the **Save** command. That is, if you use the CLI to make a change to the running-config file and then go to the Menu interface and execute a **Save** command without making a configuration change in the Menu interface, the CLI change made to the running-config file is not saved to the startup-config file. (You can still save the change by returning to the global configuration level in the CLI and executing **write memory**). For more on memory operation, see appendix C, "Switch Memory and Configuration" in the *HP ProCurve Series 2500 Switches Management and Configuration Guide* shipped with your switch and also available at <http://www.hp.com/go/hpprocurve>. (Click on technical support, then manuals.) See also "Incomplete Information on Saving Configuration Changes" on page 226.

Port Security: Changes to Retaining Learned Static Addresses Across a Reboot

Recommended Port Security Procedures

- Before configuring port security, use the switch's TFTP features to save a copy of the configuration. In the event that you later want to remove the switch's port security configuration (including MAC addresses the switch has authorized) and reconfigure port security, your task will be easier.
- If you want to manually configure the authorized MAC addresses for a port (instead of allowing the switch to learn whatever MAC addresses it detects first on the port), then prior to configuring the Static learn mode on a port, remove the LAN link from the port. This prevents the port from automatically learning MAC addresses that you do not want to include in the authorized list. After you use the **port-security <port-list> mac-address <mac-addr>** command to configure the authorized addresses you want in the list, reconnect the link.

- After you configure the authorized MAC addresses you want on a port, execute the write memory command to make these addresses permanent in the switch's configuration. (See the "Assigned/Authorized Address" bullet under "Retention of Static Addresses" in the next subsection.)

Retention of Static Addresses

Beginning with release F.02.02, port security operation has changed to the operation described below. These changes affect information provided in Table 7-1, "Port Security Parameters" on pages 7-14 and 7-15 in the *Management and Configuration Guide* (p/n 5969-2354) provided for the Series 2500 switches.

- **Learned Addresses:** In the following two cases, a port in Static learn mode retains a learned MAC address even if you subsequently reboot the switch or disable port security for that port:
 - The port learns a MAC address after you configure the port for Static learn mode in both the startup-config file and the running-config files (by executing the **write memory** command).
 - The port learns a MAC address after you configure the port for Static learn mode in only the running-config file and, after the address is learned, you execute **write memory** to configure the startup-config file to match the running-config file.

To remove an address learned using either of the preceding methods, do one of the following:

- Delete the address by using the **no port-security <port-number> mac-address <mac-addr>** command.
 - Download a previously saved configuration file that does not include the unwanted MAC address assignment.
 - Reset the switch to its factory-default configuration.
- **Assigned/Authorized Address:** If you manually assign a MAC address (using the **port-security <port-number> address-list <mac-addr>** command) and then you execute a **write memory** command, the assigned MAC address remains in memory until you do one of the following:
 - Delete it by using the **no port-security <port-number> mac-address <mac-addr>** command.
 - Download a previously saved configuration file that does not include the unwanted MAC address assignment.
 - Reset the switch to its factory-default configuration.

Disabling port security on a port does not remove an assigned MAC address from the port security configuration for that port.

Username Assignment and Prompt

Prior to release F.02.02, assigning a manager or operator username to the switch required you to use the web browser interface. Also, only the web browser interface required you to enter a username at logon if one was configured for the privilege level you were accessing. Beginning with release F.02.02 you can use the CLI **password** command to assign a manager- and/or operator-level username, and the CLI and web browser interface will require you to enter a username at logon if one is configured.

Note

On the series 2500 switches, a username is optional.

Syntax: password < manager | operator > [user-name < *user-name-str*>]

For example, to use the CLI to configure a manager user name of **sysman1** and a manager password of **top1mgr**:

```
HP2512(config)# password manager user-name sysman1
New password: *****
Please retype new password: *****
```

To use the CLI to remove all user name and password protection from the switch:

```
HP2512(config)# no password
```

Updates and Corrections for the Management and Configuration Guide

This section lists updates to the *Management and Configuration Guide* (p/n 5969-2354; August 2000).

Changes in Commands for Viewing the Current Configuration Files	page 223
Change in CLI Command for Listing Intrusion Alerts.	page 224
Changes for Listing Port and Trunk Group Statistics	page 224
Time Protocol Changes	page 224
Change in Command Line Operation	page 224
Restoring the Factory-Default Configuration	page 225
Incomplete IP Multicast (IGMP) Filtering Data	page 225
GVRP Does Not Require a Common VLAN	page 225
Incomplete Information on Saving Configuration Changes	page 226
Update to Information on Duplicate MAC Addresses Across VLANs	page 226
Incorrect Command Listing for Viewing Configuration Files	page 227
New and Corrected Information on Primary VLAN Usage.	page 227
Misleading Statement About VLANs	page 228

Changes in Commands for Viewing the Current Configuration Files

On page C-4, the manual incorrectly states that **show startup-config** displays the current startup-config file. Instead, the following is true:

- **show config:** Displays a listing of the current startup-config file.
- **show running-config:** Displays a listing of the current running-config file.
- **write terminal:** Displays a listing of the current running-config file.
- **show config status:** Compares the startup-config and running-config files and lists one of the following results:
 - Running configuration is same as the startup configuration.
This message indicates that the two configurations are the same.
 - Running configuration has been changed and needs to be saved.
This message indicates that the two configurations are different.

Change in CLI Command for Listing Intrusion Alerts

With port security configured, the switch formerly used **show interfaces** to display a port status listing that includes intrusion alerts (as described on page 7-28 in the manual). The **show interfaces** command now lists other port data (see below) and the command for listing port status with intrusion alerts is now **show interfaces brief**.

Changes for Listing Port and Trunk Group Statistics

The Port Counters screen in the Menu interface now includes flow control and broadcast limit data for each port.

The switch formerly used the **show statistics [e] <port-list>** to display port counter information (page 10-10). The command is now **show interfaces [e] <port-list>**. (The **show statistics [e] <port-list>** command is now obsolete.

Time Protocol Changes

Because the switch now offers both TimeP and SNTP (Simple Network Time Protocol) as time synchronization methods, the TimeP configuration information on pages 5-3 through 5-10 has changed. See “Enhancements in Release F.02.02” on page 153.

Change in Command Line (CLI) Operation

For the (port) Interface and VLAN commands, the command line accepts only one parameter at a time. For example, for port 1, you would use either of the following two command sets to configure duplex, flow control, and broadcast limit (instead of combining them all in one command).

At the Interface Context Level

```
HP2512(eth-1)# enable speed-duplex auto
HP2512(eth-1)# enable flow-control
HP2512(eth-1)# enable broadcast-limit 50
```

At the Global Configuration Level

```
HP2512(config)# int e 1 enable speed-duplex auto
HP2512(config)# int e 1 enable flow-control
HP2512(config)# int e 1 enable broadcast-limit 50
```

This change affects the following commands:

Interface Commands	VLAN Commands
broadcast-limit	forbid
disable	tagged
enable	untagged
flow-control	
lacp	
monitor	
speed-duplex	
unknown-vlans	

Restoring the Factory-Default Configuration, Including Usernames and Passwords

Page 11-20 in the Management and Configuration guide incorrectly implies that the **erase startup-config** command clears passwords. This command does reset the switch to its factory-default configuration, *but does not remove any user names or passwords (Manager or Operator) configured in the switch.* To remove user names and passwords, do any one of the following:

- Execute the **no password** command in the CLI.
- Select the **Delete Password Protection** option in the "Set Password" menu screen.
- Press and hold the Clear button on the switch for one second.
- Restore the factory-default configuration by using the Clear/Reset button combination, as described under "Restoring the Factory Default Configuration" in the "Troubleshooting" chapter of the *Installation and Getting Started Guide* you received with the switch.

Incomplete IP Multicast (IGMP) Filtering Data

The Note on page 9-92 in the *Management and Configuration Guide* states that "IGMP requires an IP address and subnet mask for any VLAN used for IGMP traffic." This is no longer true. See "Enhancements in Release F.02.02" on page 135.

The second paragraph in the note on page 9-101 in the *Management and Configuration Guide* provides incomplete data on the "well-known" or reserved IP multicast addresses that IGMP does not filter in the Series 2500 switches. See "The Switch Excludes Well-Known or Reserved Multicast Addresses from IP Multicast Filtering" on page 219.

GVRP Does Not Require a Common VLAN

Delete the note at the top of page 9-78 in the *Management and Configuration Guide*. GVRP does not require a common VLAN (VID) connecting all of the GVRP-aware devices in the network to carry GVRP packets.

Incomplete Information on Saving Configuration Changes

Using the CLI to make a configuration change to the running-config file, then going to the Menu interface and making another configuration change, and then executing the Menu interface **Save** command saves all of your changes to the startup-config file. (At this point, the startup-config file and the running-config file will have identical configurations, and will contain all of the changes that you made in both interfaces.)

The second paragraph of the Note on page C-6 in the *Management and Configuration Guide* states that "Using the Save command in the menu interface will not save a change made to the running config by the CLI." This statement is true where you:

1. Make configuration changes in the CLI
2. Move to the Menu interface, but make no configuration changes while using the Menu interface.
3. Execute the **Save** command in a Menu interface screen.

However, the statement is not true if you make a configuration change in the Menu interface before going to step 3, above. See also "Switch Memory Operation" on page 220.

Update to Information on Duplicate MAC Addresses Across VLANs

On page 9-75 of the *Management and Configuration Guide*, the following information replaces the text in the fourth bullet from the top and the Note:

Duplicate MAC addresses on different VLANs are not supported and can cause VLAN operating problems. These duplicates are possible and common in situations involving Sun workstations with multiple network interface cards, with DECnet routers, the ProCurve routing switches (9304M, 9308M, and 6308M-SX), and with certain Hewlett-Packard routers using OS versions earlier than A.09.70 where any of the following are enabled: IPX, IP Host-Only, STP, XNS, DECnet, and possibly others. When in doubt, ask your router vendor under what conditions, if any, the router uses the same MAC address on more than one interface. Regarding the HP ProCurve routing switches, see the FAQ "Q: What is the recommended way to connect multiple VLANs between a routing switch and a layer 2 switch?" on page 253.

Note

Duplicate MAC addresses are likely to occur in VLAN environments where XNS and DECnet are used. For this reason, using VLANs in XNS and DECnet environments is not currently supported.

On page 11-10 of the *Management and Configuration Guide*, under "Duplicate MAC Addresses Across VLANs", the text suggests that duplicate MAC addresses on separate VLANs can cause VLAN operating problems. However, duplicate MAC addresses on different VLANs may cause operating problems that have no apparent connection to VLAN operation. Thus, in the paragraph under "Duplicate MAC Addresses Across VLANs", delete the word "VLAN" from the first sentence. That is, the sentence should be: "Duplicate MAC addresses on different VLANs are not supported and can cause operating problems."

Incorrect Command Listing for Viewing Configuration Files

On page C-4 of the *Management and Configuration Guide*, under "How To Use the CLI To View the Current Configuration Files", the **show startup config** command is incorrect. Use the following "show" methods for listing configuration files:

- **show config** : Displays the startup-config file.
- **show config run** : Displays the running-config file.

(The **write terminal** command also displays the running-config file.)

The **show config**, **show config run**, and **write terminal** commands list the following configuration data:

- Daylight Time Rule setting
- Hostname (system name)
- SNMP server community name and status
- The default VLAN and its IP address setting
- Any other configuration settings that differ from the switch's factory-default configuration.

New and Corrected Information on Primary VLAN Usage

The second bulleted item on page 9-54 incorrectly states that "The switch reads DHCP responses on the primary VLAN instead of on the default VLAN." The switch reads DHCP (and Bootp) responses received on all VLANs. The restriction is that the switch only honors default gateway addresses, TimeP server addresses, and IP TTL values learned from DHCP or Bootp packets received on the primary VLAN.

Also on page 9-54, add the following item to the bulleted list:

- When TimeP is enabled and configured for DHCP operation, the switch learns of TimeP servers from DHCP and Bootp packets received on the primary VLAN.

Misleading Statement About VLANs

On page 9-56 in the Management and Configuration Guide, the last sentence in item 1 implies that by default the switch is configured for eight VLANs. The sentence should read as follows:

"By default, VLAN support is enabled to support up to eight VLANs, and the switch is configured for one VLAN (the default VLAN). By changing the Maximum VLANs to support parameter, you can configure up to 29 VLANs."

Software Fixes

Release F.01.07 was the first software release for the HP ProCurve Series 2500 switches.

Release F.01.08	page 230
Release F.01.09 (Beta Release Only)	page 230
Release F.01.10	page 230
Release F.02.02	page 230
Release F.02.03	page 232
Release F.02.04 (Beta Release Only)	page 233
Release F.02.05 (Beta Release Only)	page 234
Release F.02.06 (Beta Release Only)	page 235
Release F.02.07 (Beta Release Only)	page 235
Release F.02.08 (Beta Release Only)	page 236
Release F.02.09	page 236
Release F.02.10	page 236
Release F.02.11	page 236
Release F.02.12	page 237
Release F.02.13	page 237
Release F.04.01	page 237
Release F.04.02	page 238
Release F.04.03	page 238
Release F.04.04	page 238
Release F.04.08	page 239
Release F.04.09 (Beta Release Only)	page 239
Release F.05.05 (Beta Release Only)	page 239
Release F.05.09 (Beta Release Only)	page 243
Release F.05.10 (Beta Release Only)	page 243
Release F.05.12 (Beta Release Only)	page 243
Release F.05.13 (Beta Release Only)	page 244
Release F.05.15 (Beta Release Only)	page 244

Software Fixes

Release F.05.16 (Beta Release Only) page 246

Release F.05.17..... page 247

Release F.01.08

Fixed in release F.01.08:

- **100/1000-T transceiver** — When using this 100/1000-T transceiver and negotiating to 100 Mbps, the port may report that it is operating at 100 full duplex, when it is actually operating at 100 half duplex.
- **Web-Browser Interface** — The product label in the web-browser display for the Switch 2512 is incorrectly displayed as **Switch 2524**.

Release F.01.09 (Beta Release Only)

Fixed in release F.01.09:

- **Console/Management** — A console and management (SNMP, telnet, etc.) hang may occur when an illegal MAC address is detected on a port configured with a port security action of "send-disable".

Release F.01.10

Fixed in release F.01.10:

- **Port Security** — The switch does not send an alarm upon a port security violation when the port security learn-mode is "continuous" and the action is "send-alarm".
- **Port Security** — If the configuration is not saved (i.e., **write mem**) before the switch is rebooted, the learned addresses are not saved.
- **Port Security** — A port that has been disabled due to a security violation does not remain disabled after the switch is rebooted.

Release F.02.02

This release adds the following new features:

- TACACS+
- CDP (Cisco® Discovery Protocol)
- SNMP
- Improved IGMP capabilities

For details about the above enhancements, refer to “Enhancements in Release F.02.02” on page -153.

Note

The startup-config file saved under version F.02.02 is NOT backward-compatible with previous software versions. HP recommends that you save a copy of the pre-02.02 startup-config file BEFORE UPGRADING to F.02.02 or greater, in case there is ever a need to revert back to pre-02.02 software. Instructions for saving a copy of the startup-config file are found in the "Transferring Switch Configurations" section of Appendix A in the *Management and Configuration Guide* available for the switch.

Fixed in release F.02.02:

- **100/1000-T Transceiver** — After switch is rebooted, the port counters contain an incorrect large value.
- **100/1000-T Transceiver** — If the transceiver speed-senses from 1000 Mbps to 100 Mbps (or the reverse), the port incorrectly stays at the previous speed (i.e., speed mismatch) while the switch incorrectly shows linkbeat for that port. [Fix is to have the switch not establish linkbeat. The switch must be rebooted in order to establish linkbeat after the transceiver speed-senses from 1000 Mbps to 100 Mbps or vice versa.]
- **ARP** — If switch's gateway is the same as its own IP address, switch cannot ping off-net and "show arp" output does not include gateway, after pinging the configured gateway.
- **CLI** — The output of the **show help** command from the operator level context lists commands that are unavailable due to insufficient privileges and the output of the **show interface config** command does not properly align the trunk designations within the **Port** column.
- **Config** — When a config is reloaded that was saved off from a switch, it does not match the config offloaded as follows:
 - a. SNMP community parameter **unrestricted** is changed to **(null)**.
 - b. **forbid** commands are added to the VLAN configuration.
- **Console** — If an active port is configured as disabled and, while the port is disabled its trunk membership is changed, the switch console becomes inaccessible.
- **Fault-Finder** — The fault-finder configuration as reported by **show config** or **write term** does not correctly display the type of alarm.
- **IP** — The **IP Config** parameter changes from **DHCP/Bootp** to **Manual** on the default VLAN when trying to add a new VLAN address that is the same (i.e., duplicate) as the DHCP-acquired IP address of the default VLAN. [Fix is: error message is generated when the user attempts to configure a duplicate IP address.]

Software Fixes

- **LACP** — Resolves several issues with LACP, including: conversation on a trunk may momentarily fail if a trunk member port goes down, difficulty accessing the MIB, configuration issues, port priority issues, problems with dynamic negotiation, and switch crashes with messages similar to:

```
-> Software Exception at woody_dev.c: 450 in AdMgrCtrl
-> ppmgr_setDefaultPriority: invalid port number
```

and

```
-> Software exception at woody_pktDriver.c:317 -- in 'eDrvPoll'
-> ERROR: ASIC buffer return failure
```

- **Link** — The switch exhibits intermittent link behavior when connected to some 3C905B 3Com NICs.
- **Monitor Port** — If a user attempts to monitor the monitoring port the switch displays a meaningless error message.
- **Ping** — The switch replies to pings with a source address of 127.1.1.1, which is a loopback address.
- **Port Security** — Static addresses are saved to startup-config without the user executing a **write memory** command. [Fix is: static addresses will only be saved to startup-config by executing the **write memory** command.]
- **SNMP** — For ports with no transceiver present, any SNMP sets to the **hpSwitchPortTable** fail and an SNMP get of **hpSwitchPortType**, **hpSwitchStpPortType**, or **hpSwitchIcmpPortType2** returns an illegal value of zero (0).
- **Stack Management** — Resolves several issues with ProCurve Stack Management via the web-browser interface, including problems with stacking configuration screen, Stack Member port counters, and not being able to add a candidate to a stack.
- **STP** — Resolves several issues with STP, including problems with an SNMP set and get of the **dot1dStpPortEnable** MIB variable, setting STP parameters via SNMP disables Spanning Tree, and a switch crash with a message similar to:

```
-> Software exception at stp_ctrl.c:154 -- in 'mStpCtrl'
```
- **TFTP/XMODEM** — The switch's event log is not properly formatted when captured via TFTP or XMODEM using the **copy** command.
- **VLAN** — After creating several VLANs, the default value for all ports in one VLAN is **forbid** and this value cannot be changed.

Release F.02.03

Fixed in release F.02.03:

- **Stack Management** — Cannot access member switches via SNMPv2c.

Release F.02.04 (Beta Release Only)

The switch's CDP packets have been modified to better interoperate with older Cisco IOS versions. Certain legal CDP packets sent from the ProCurve switch could result in Cisco routers, running older IOS versions, to crash.

Note

The ProCurve switch's CDP packets are legal both before and after this modification.

Fixed in release F.02.04:

- **Buffer Leak** — A message buffer leak occurs when the switch receives a TACACS+ 'DISC' character.
- **CDP** — The switch sends the wrong MAC address for itself in CDP packets.
- **Console/TELNET** — The switch console may hang, or TELNET session may become inaccessible, if either of the following conditions occur:
 - While using TELNET, if the inactivity timer ends the session, subsequent attempts to re-establish the TELNET session may result in the user's login failing at the login prompt.
 - If a console session is ended due to inactivity timer expiration, the user is not able to establish another console session.
- **Continuous Reboot** — The switch continuously reboots upon downloading a configuration file containing a IP configuration (from DHCP or BootP).
- **Crash** — The switch may crash with a message similar to:


```
-> Software exception at infTrunks.c:264 in 'mAdMgrCtrl'.
```

 This crash may occur if both the following conditions exist:
 - All ports of a dynamic trunk are off-line (for example, disconnected); and
 - The trunk is a member of the default VLAN.
- **Crash** — At very high levels of traffic, the switch may crash with a message similar to:


```
-> Software exception at xcvr_util.c:1387 -- in 'mPmSlvCtrl'
```
- **DHCP** — The DHCP address of the TimeP Server is not displayed in the output of the **show ip** CLI command or in the IP configuration menu screen.
- **IGMP** — If there are several IGMP groups in several VLANs, the switch may delete IGMP multicast groups from its table, resulting in flooded multicasts.

Software Fixes

- **IGMP** — If there are several IGMP groups in several VLANs, and the switch is acting as Querier, the switch may stop sending IGMP Queries on some of its VLANs.
- **IGMP** — All Querier intervals on the switch will be cut in half if IGMP, after already being enabled, is disabled and then re-enabled.
- **IGMP** — The switch does not fully support 256 IGMP groups, as intended. For example, with 15 VLANs and 40 IGMP groups, the 40th group gets flooded.
- **LED** — The MAX mode LED does not turn on for port where Gigabit Stacking Transceiver is installed.
- **Memory Leak and Crash** — If the "Send Authentication Traps" trap receiver parameter on a Member switch is set to "Yes", it will cause a memory leak on the Stack Commander switch. The memory leak can eventually cause a crash. The specific details of the crash vary.
- **Port security** — Port security learn mode and the learned MAC addresses are not saved after the switch is rebooted.
- **Port Security** — With port security on, the switch does not remember learned static MAC addresses after reboot.
- **Stack Management** — The commander may hang (SNMP, ping, TELNET, etc.) and other CPU functions may stop when the switch is queried by management applications such as the WhatsUp Gold utility.
- **Stack Management** — If a commander has a CDP neighbor, the commander may run out of packet buffers and hang (SNMP, ping, TELNET, etc.).
- **TELNET** — If a TELNET session times out due to the inactivity timer expiring, then a subsequent TELNET session will freeze at the switch's copyright screen, before displaying "Press any key to continue". Note: This does not affect console sessions.
- **TELNET** — Unable to open new TELNET sessions due to switch not correctly closing previous TELNET sessions.
- **Web-browser interface** — Clicking the stack management close-up button does not show the 4108GL switch.

Release F.02.05 (Beta Release Only)

Added new Isolated Port Groups feature. Each switch port is configurable as any one of four types:

- Public
- Private
- Local
- Uplink

Note

Contact your local Customer Care Center before activating this feature to receive proper configuration instructions. Failure to configure this feature properly will result in unexpected connectivity problems.

Release F.02.06 (Beta Release Only)

Textual modifications made to the Isolated Port Groups feature.

Release F.02.07 (Beta Release Only)

This release adds two new features:

- Spanning Tree fast "uplink" mode
- **show tech** command (Captures information to help with troubleshooting.)

The above features are available on HP's ProCurve website in release F.02.11. For more information, turn to "Enhancements in Release F.02.11" on page 119.

Fixed in F.02.07:

- **Bus Error** — The switch may crash with a bus error if its IP address is changed during a telnet session (originated from the switch).
- **Crash** — If the switch's DHCP-learned IP address is a duplicate with another node's IP address, the switch may crash with a message similar to:

```
-> software exception at alloc_free.c:432 -- in 'eDrvPoll'
-> buf_free: corrupted buffer.
```

- **Performance** — Slow performance and possible packet loss when switch was connected to Intel 10/100 NICs.
- **Performance** — Slow performance over 10 Mbit half-duplex links when switch is connected to various NICs such as 3COM 3C905B, 3COM 3C590C, D-Link DE-528, and Lantech PCI-NET/32T.
- **Transceiver hot-swap** — A transceiver hot-swap is falsely reported when the screws on a transceiver are tightened or loosened. The event log will report a message similar to:

```
I 01/01/90 00:00:19 ports: port 13: Xcvr Hot-Swap detected. Need
reboot.
```

Software Fixes

- **XRMON** — Various XRMON counters display incorrect values. Possible symptoms include network management applications reporting a too high network utilization (TopTools may report "crossed octets").

Release F.02.08 (Beta Release Only)

Fixed in F.02.08:

- **Crash** — If a transceiver is repeatedly installed and removed, the switch may crash with a message similar to:

```
-> Software exception at woodyDma_recv.c:154 -- in 'eDrvPoll'
```

Release F.02.09

Fixed in F.02.09:

- **Configuration download** — Downloading a configuration file (via TFTP or Xmodem) sometimes failed to reboot the switch.
- **Isolated Port Groups** — Downloading a configuration file (via TFTP or Xmodem) containing port isolation commands may fail with error messages similar to:

```
line: 6. Error setting configuration tree.  
Corrupted download file.
```

Release F.02.10

Fixed in release F.02.10:

- **LEDs/Port toggling** — The switch LEDs flash randomly on various ports (even ports that do not have cables attached) when a 100/1000-T transceiver is installed. Excessive port toggling may also occur on ports that have cables attached. These problems have been associated with network management applications such as TopTools.

Release F.02.11

Fixed in release F.02.11

- **Auto-TFTP** — If the switch's configuration file contains STP (i.e., STP is enabled), auto-tftp does not download a new OS.
- **Transceivers** — Removing and re-inserting both transceivers simultaneously many times with network cables attached and without an intervening reboot may cause the switch to crash with a message similar to:

```
-> Software exception in ISR at buffers.c:1627
```

Release F.02.12

Fixed in release F.02.12

- **Monitoring Port** — When a config file containing a Monitoring Port configuration is loaded onto the switch via TFTP or XModem, the Monitoring Port feature does not work properly.

Release F.02.13

Fixed in release F.02.13

- **Monitoring Port** — Monitoring Port configuration changes made within a particular switch interface (e.g., web-browser interface), are not correctly displayed within the other switch interfaces (e.g., CLI and Menu).

Release F.04.01 (Beta Release Only)

Fixed in release F.04.01

- **CLI** — The response to an incomplete trunk configuration command did not produce the proper message "Incomplete input: Trunk."
- **CLI** — The crash history is lost after the "reload" command is performed from the CLI.
- **Crash** — A transceiver hot-swap may cause the switch to crash with a message similar to:

```
-> Software exception at woodyDma_rev.c:154 -- in 'eDrv'.
```
- **Crash** — A transceiver hot-swap may cause the switch to crash with a message similar to:

```
-> Software exception in ISR at buffers.c:1627.
```
- **Crash** — The switch may crash with a message similar to:

```
-> Software exception at woodyDma_recv.c:154 -- in 'eDrv'.
```

This crash may occur if both the following conditions exist:

 - a. The "reload" CLI command is issued; and
 - b. A 100/1000-T transceiver is installed
- **Flow Control** — Changing Flow Control setting on a port is not reflected in Auto-negotiation's advertised capability.
- **IGMP** — Interoperability issues with some Cisco devices cause IGMP groups to be aged out of the switch's IGMP tables prematurely.
- **Menu/Web-Browser Interface** — Display of mirror port configuration is inconsistent between menu and WEB interface.

Software Fixes

- **Port Configuration** — Changing a port setting from one Auto mode to another may not be reflected in Auto-negotiation's advertised capability without a switch reset, or module hot-swap.
- **Port Monitoring** — Port monitoring does not work correctly after a TFTP transfer of the configuration from the switch to the server and then back to the switch.
- **Stack Management** — Master switch was not properly making security checks when passing information along to a member switch.
- **TFTP** — Menu and browser displays of switch configuration are not accurate after a TFTP transfer of the switch config file to the switch. Only occurs when a port is configured for network monitoring.
- **VARIOUS: Crash/Bus Error** — A Get request of a specific long OID can result in a bus error, an agent hang, or a switch crash with a message similar to:

```
-> Software_exception at svc_misc.s:379 -- in mCdpCtrl  
malloc_else_fatal() ran out of memory
```
- **Web-Browser Interface** — Web display of port utility window did not display port H24.
- **Web-Browser Interface** — User could input an invalid MAC address, i.e. multicast or broadcast address, in the security policy field.
- **Web-Browser Interface** — Incorrect font size used in VLAN configuration screen.

Release F.04.02 (Beta Release Only)

Fixed in release F.04.02

- **Corrupted Flash** — An SNMP set, during the OS download operation of TopTools, while the switch is writing new OS to flash may result in corrupted flash and switch may boot up in LAN Monitor mode.

Release F.04.03 (Beta Release Only)

Fixed in release F.04.03

Modification of Lab troubleshooting commands.

Release F.04.04 (Beta Release Only)

Fixed in release F.04.04

Modification of Lab troubleshooting commands.

Release F.04.08

Fixed in release F.04.08

Modification of Lab troubleshooting commands.

Release F.04.09 (Beta Release Only)

Fixed in release F.04.09

- **Agent Hang** — Agent processes (such as console, telnet, STP, ping, etc.) may stop functioning when the IGMP querier function is disabled, and then re-enabled, on a VLAN that does not have an IP address configured.
- **Agent Hang** — Agent processes (such as console, telnet, STP, ping, etc.) may stop functioning. This agent hang has been associated with the CERT SNMPv1 "encoding" test #1150.
- **Agent Hang** — Agent processes (such as console, telnet, STP, ping, etc.) may stop functioning. This agent hang has been associated with the X2 SSH utility.
- **CLI** — When reaching the inactivity timeout expiration after typing the CLI command "enable" in a telnet session at operator mode, the text in the CLI prompt may get corrupted with text similar to `gfs_alp_104Null Varbind`.
- **Crash** — When hot-swapping transceivers multiple times, the switch may crash with a message similar to:

```
-> Software exception at port_sm.c:378 in -- 'mPmSlvCtrl'
```
- **STP/RSTP** — Port path cost is reset even though path cost is configured for "Auto".

Release F.05.05 (Beta Release Only)

Time Zone Issue

Starting with the F.05.*xx* version of the switch operating system software, the method of configuring the Time Zone for TimeP or SNTP configuration has been updated. Previous switch software, for all HP ProCurve switches, used positive time offset values for time zones that are West of GMT and negative values for time zones that are East of GMT. The standards indicate that time zones West of GMT should be designated by negative offset values, and time zones East of GMT by positive values. Software version F.05.*xx* updates this configuration method, but if you use the same values for indicating time zones as you did for previous HP ProCurve switches, the time will be set incorrectly on your HP ProCurve Switches 2512 and 2524. For example, for previous HP ProCurve switches, the US Pacific time zone was configured by entering +480. With software version F.05.*xx*, the US Pacific time zone must now be configured by entering -480.

Note

The startup-config file saved under version F.05.05, or later, is NOT backward-compatible with previous software versions. The user is advised to save a copy of the pre-05.05 startup-config file BEFORE UPGRADING to F.05.05 or greater, in case there is ever a need to revert back to pre-05.05 software. Instructions for saving a copy of the startup-config file are found in the "Transferring Switch Configurations" section of Appendix A in the *Management and Configuration Guide* (included on the Product Documentation CD-ROM (PDF format)) that shipped with the switch.

Fixed in release F.05.05

- **Agent Hang** — Agent processes (such as console, telnet, STP, ping, etc.) may stop functioning when the IGMP querier function is disabled, and then re-enabled, on a VLAN that does not have an IP address configured.
- **ARP** — Changing the IP address of a VLAN does not delete the ARP entry for the old IP address.
- **GARP/Event log** — Garp event log messages may be garbled.
- **CLI** — The CLI command **show arp** displays the wrong port number for some ARP entries.
- **CLI** — The CLI command **show trunks** lists incorrect information for dynamic trunks.
- **CLI** — The CLI command **getmib** with no parameters returns the message `Incomplete input: - EOI -`.
- **CLI** — Unrelated information was shown at the end of the CLI command "show vlan 1" output.
- **CLI** — Command "no qos" did not reset port priority to "0".
- **CLI** — When reaching the inactivity timeout expiration after typing the CLI command "enable" in a telnet session at operator mode, the text in the CLI prompt may get corrupted with text similar to "gfs_alp_104Null Varbind".
- **CLI** — The CLI command "show tech" causes an error message when the command is executed from within config mode.
- **CLI** — The prompt for saving the config does not handle a DISC character appropriately.
- **CLI/Timezone** — The switch time is wrong if CLI used to set timezone and timezone may not operate properly after switch is rebooted. West of GMT is now a negative offset and east of GMT is now a positive offset.

- **Crash** — If dynamic trunks are configured and the switch is rebooted, the switch may crash with a message similar to:

```
->Software exception at rstp_dyn_reconfig.c:243 in -- 'Lpmgr'
```

- **Crash** — The "show config" CLI command may cause the switch to crash with a message similar to:

```
->Software exception "xlate.c:1358 in 'mSess1'
```

- **Crash** — When hot-swapping transceivers multiple times, the switch may crash with a message similar to:

```
-> Software exception at port_sm.c:378 in -- 'mPmSlvCtrl'
```

- **Crash** — The switch may crash with a message similar to:

```
-> Software exception at alloc_free.c:545 -- in 'eDrvPoll'
```

- **Crash** — The switch may crash with a message similar to:

```
->Asserts in rv.cc line 632
```

- **Crash** — When adding ports to a manual LACP trunk configuration, the switch may crash with a message similar to:

```
-> Assert line 10070 cli_config_action.c
```

- **Crash** — The switch may crash with a message similar to:

```
-> Bus Error: HW Addr=0x0000000 IP=0x002fe640 Task='eTelnetd'
```

This crash has been associated with security/vulnerability test applications such as Nessus.

- **Event Log** — Log messages for trunks and trunk members enhanced to be easier to read.
- **GVRP** — Upstream GVRP neighbor does not pass VLAN information to downstream neighbor after a topology change.
- **IGMP** — Interoperability issues with some Cisco devices (such as some Cisco Catalyst 5000 & 6000 series switches) cause IGMP groups to be aged out of the switch's IGMP tables prematurely.
- **LACP/802.1x** — 802.1x and LACP trunks can co-exist on the same port. (Fix is to make these trunks mutually exclusive.)
- **LACP** — LACP maintains a dynamic trunk with only 1 port configured for the trunk group.
- **Link-up polling interval** — A delay of up to 1.7 seconds between plugging in a cable (linkbeat established) and traffic being forwarded to and from that port may cause problems with some time sensitive applications. For example, AppleTalk dynamic address negotiation can be affected, resulting in multiple devices using the same AppleTalk address.

Software Fixes

- **Loop/VTP** — The switch will incorrectly forward VTP packets from third party devices if that packet is received on a blocked port.
- **Menu** — The switch may not display the complete forwarding table while performing a MAC address search.
- **Menu** — Incorrect error message is displayed when attempting to create a VLAN which exceeds the maximum number of VLANs supported. (Fix is to display an error message similar to "VLAN limit already reached.")
- **Menu** — Menu does not allow a port configuration change from a full-duplex/flow control setting to a half-duplex/no flow control setting.
- **Menu** — Menu does not allow trunks to be configured on transceiver ports that have never had a transceiver installed before.
- **Menu/CLI** — Modified help message for RSTP.
- **Menu/VLAN** — The VLAN help text has been modified
- **NNM/Stacking** — If stacking is configured, NNM cannot discover the device as a generic switch.
- **Performance/Crash** — Slow performance may occur when using 10/100 ports or the 100FX transceiver operating at half-duplex. And when using 100FX, Gigabit Stacking, Gigabit-SX, or Gigabit-LX transceivers operating at full-duplex (Note: The Gigabit transceivers can only operate in full-duplex mode.) This was due to the InterPacket Gap (IPG) being too long for half-duplex and too short for full-duplex. This may also result in a switch crash with a message similar to:

```
-> Software exception at woodyDma_recv.c:154 -- in 'eDrvPollRx'  
-> DMA unrecoverable error
```
- **STP** — Switch does not forward STP BPDUs when STP is disabled.
- **STP Fast Mode** — A port configured for STP fast mode behaves like a standard STP port.
- **STP/RSTP** — Port path cost is reset even though path cost is configured for "Auto".
- **STP/Running-Config** — STP path-cost is not written to the configuration when using the CLI.
- **STP/Startup-Config** — When a startup-config file containing an 802.1d STP configuration is reloaded that was saved off from the switch, an error similar to the following occurs:

```
Line: 13. Invalid input: stp802.1d  
Corrupted download file.
```

- **TACACS+** — When logging into the switch via TACACS+ encrypted authentication, the packet header has the 'encryption' field set to 'TAC_PLUS_CLEAR' when the body of the packet is actually encrypted.
- **Time Zone** — The time zone changes made for PR_4524 are not transparent to the end user.
- **VLAN** — New vlans have no members when creating/deleting lots (>30) of unique vids.
- **VTP/ISL** — The switch does not forward Cisco VLAN Trunk Protocol (VTP) or Inter-Switch Link (ISL) packets.
- **Web-Browser Interface** — After clearing the intrusion flag in the web-browser interface, the intruder flags are not removed.
- **Web-Browser Interface** — Configuring static learning for port security may result in the error message "error in pdu".
- **Web-Browser Interface** — The product registration screen contains a typographical error. (Fix: The phrase "...does not appears above..." is now "...does not appear above...".)

Release F.05.09 (Beta Release Only)

Fixed in release F.05.09

- **GVRP** — If GVRP is disabled on a port (to prevent GVRP from being active on that port), then any VLANs statically configured on that port will not be advertised out (other) GVRP-active ports.

Release F.05.10 (Beta Release Only)

Fixed in release F.05.10

- **STP** — The issue only occurs when this device, running STP in 802.1w mode, has an 802.1d immediate upstream neighbor. Various possible symptoms, including: i. Topology Change count increments an excessive number of times after a physical topology change. (Note: Even with the fix, the Topology Change count increments by more than one for a single change, which is expected behavior.); ii. Topology Change count increments for much longer than 30 seconds after a physical topology change; iii. After a physical topology change, the spanning tree may take a long time to re-converge, and may never re-converge; iv. Possible flooding storms (which users may mistakenly report as broadcast storms).

Release F.05.12 (Beta Release Only)

Adds the following enhancement:

- Changes to 802.1x to support Open VLAN Mode

Software Fixes

Release F.05.13 (Beta Release Only)

Adds the following enhancement:

- Changes to Isolated Port Groups to add two new groups: group1 and group2.

Release F.05.15 (Beta Release Only)

Adds the following enhancements:

- Increased IGMP V3 interoperability by allowing the switch to keep (and not prune) V3 groups. This lets the switch interoperate in an IGMP V3 environment without pruning off the V3 groups (due to the Data-Driven IGMP feature) or always flooding.
- Display the IGMP Querier in the output of the CLI **show ip igmp** command.

Fixed in release F.05.15

- **Bcast limit** — Configuring broadcast limiting does not take effect until a reboot.
- **CDP** — When CDP is disabled, the switch does not forward CDP packets.
- **Configuration file** — In the configuration file, the command:

```
port-security 1 learn-mode continuous action send-alarm
```

results in this error message:

```
Inconsistent Value
```
- **Configuration file** — In the configuration file, a command, such as:

```
port-security 1 learn-mode continuous action send-alarm
```

results in this error message: `Inconsistent Value`.
- **Counters** — Counters reset when you hot-swap a Gig-T transceiver.
- **Counters** — The Switch does not distinguish between CRC and alignment errors.
- **Crash** — When setting the host name to a very long (~20 characters) string, the switch may crash with a bus error similar to:

```
-> Bus error: HW Addr=0x29283030 IP=0x002086ac Task='mSnmprCtrl'  
  
Task ID=0x165ae00.
```
- **Flow control** — Users are allowed to configure flow control for half-duplex ports, even though the switch does not support flow control ("back pressure") for half-duplex links.
- **Flow control** — Users are allowed to configure 802.3x flow control for half-duplex ports.

- **IGMP** — Checking whether an IP DA and/or an IGMP Group Address is a valid IP multicast address before taking any IGMP action on it.
- **IGMP** — Fixed Group-Specific Query (GSQ) timing in Normal Leave case to be a minimum of 1 second (as IGMP standard specifies and as the GSQs advertise). This occurs when the Querier forces an interval between GSQs. The internal GSQ timer value was increased to force this change; otherwise, the time between when the timer is set and when it is serviced can result in GSQs that are usually .2-sec to .6-sec apart, instead of the minimum 1-sec. Prior to the fix, the switch sometimes pruned a group only ~3/4 of a second after sending out a GSQ. Some end nodes take slightly longer than this to reply to the GSQ.
- **IGMP** — Fixed the case where IGMPv3 Join contains an invalid IP Mcast address or a reserved IP Mcast address in the IGMP Group Address field. Previously, the switch would attempt to stop processing the Join and mistakenly double-free (one symptom can be a "software exception at alloc_free.c ... buf_free: corrupted buffer") or double-forward the Join packet. Now, the switch will simply bypass processing the "offending" sub-record and continue with the rest of the Join, freeing or forwarding the packet only once.
- **IGMP** — Not currently checking whether an IP DA and/or an IGMP Group Address is a valid IP multicast address before taking any IGMP action on it.
- **IGMP** — A Group-Specific Query (GSQ) timeout is currently .2 to .6 seconds, rather than the specified default of 1 second.
- **IGMP** — When an IGMP v3 Join contains an invalid IP Multicast address or a reserved IP Multicast address in the IGMP Group Address field, the switch will attempt to stop processing the Join, and mistakenly double-free, or double-forward the Join packet. One possible symptom is a switch crash similar to:


```
software exception at alloc_free.c ... buf_free: corrupted buffer
```
- **Memory Leak** — Related to the "WhatsUp Gold" network management application. Triggered when configuring the enable/login password on the switch.
- **Security/Vulnerability** — "Cross-site scripting" issue. One of the Nessus (see www.nessus.org) tests fails, reporting: "Vulnerability found on port http (80/tcp).
- **SNMP** — The OID ifAlias is defaulted to "not assigned", causing Network Node Manager to log error messages. (The fix is to default ifAlias to a zero-length string, as stated in the MIB, or make each port have a unique value.)
- **SNMP** — The switch does not support community names other than PUBLIC in traps.
- **SNMP/Crash** — A walkmib of the cdpCacheDeviceId OID using an ifIndex value of the Default VLAN causes the switch to crash with a bus error similar to:


```
-> Bus error: HW Addr=0x5265766d IP=0x002592e8 Task='mSnmprCtrl'
      Task ID=0x12c2158 fp: 0x00000005 sp:0x012c1e28 lr:0x00259430
```

Software Fixes

- **STP** — Under some conditions, an 802.1w non-Root switch will have a zero Root Path Cost.
- **TACACS+** — The TACACS server IP is shown on the 'splash screen'.
- **TELNET** — TCP port 1506 is always open.
- **UI** — In the absence of a time server, the switch may report that it is the year "26".
- **Web** — Web browser port utilization pop-up does not display the bandwidth number. Shows x% of 0Mb instead of x% of 100Mb or x% of 1Gb.
- **Web/Security** — When a 2500 series switch is acting as an IP Stack Commander, it is possible to send a specific command to the 2500's WEB agent that will inappropriately reset an IP Stack Member switch.
- **Web/Security** —
 1. Set Series 2500 switch as commander of a stack containing only one 1600M/2400M/2424M/4000M/8000M switch member.
 2. Send the following URL to the Series 2500 switch (Commander):
`http://<IP ADDRESS>/sw2/cgi/device_reset?`This will cause the member switch to reboot, even without a password on the commander (password not allowed on a member).

Release F.05.16 (Beta Release Only)

Fixed in release F.05.16

- **IP Stack Mgmt/Web** — A bus error occurs when accessing the close-up view of a 15-member stack (IP Stack Management) through the Web interface.
- **LACP** — With a 2500-Series switch linked via SX or LX to 5300-Series switch, turn off LACP on the 2500 when RSTP is enabled, and STP's status becomes disabled on that port.
- **LACP** — With a 2500-Series switch linked via SX or LX to a 5300-Series switch, turn off LACP on the 2500 when RSTP is enabled, and STP's status becomes disabled on that port.
- **RSTP/LACP** — Turning LACP off, then back on, leaves LACP in Passive mode. This can result in loops (broadcast storms, etc.)

- **Trunking** — With ports 25 and 26 configured in a trunk group, the **show trunk 25,26** command displays incorrect information for Trunk Group Name and Trunk Group Type.

Example output:

Port	Name	Type	Group	Type
25		1000SX	Trk1	Trunk
26		1000SX	1000SX	1000SX

- **Web** — Sun java v1.3.x and v1.4.x interoperability issue: high CPU utilization.
- **Web** — Sun java v1.3.x and v1.4.x interoperability issue resulting in high CPU utilization on the switch.
- **Web/Stack Mgmt** — Software version isn't displayed in web-agent identity screen.
- **Web/Stack Mgmt** — Inverted IP address displayed in the Identity tab when the member is accessed through the commander.
- **Web/Stack Mgmt** — Inverted IP address displayed in the Identity tab when the IP Stack Member switch is accessed through the IP Stack Commander switch.

Release F.05.16

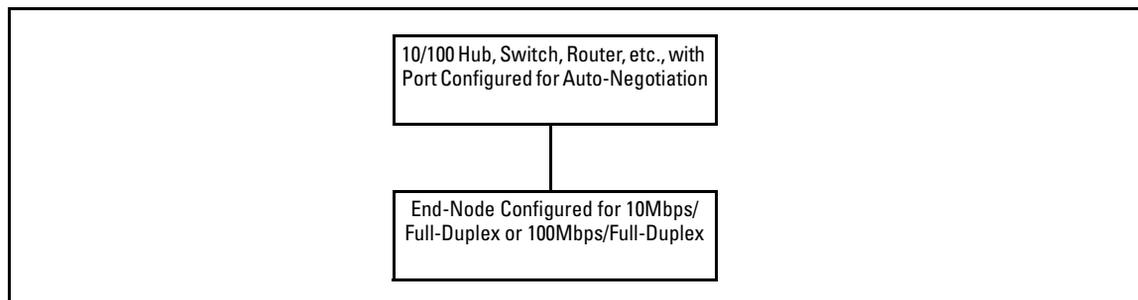
Modification of Manufacturing test commands.

Series 2500 FAQs from the HP ProCurve Website

Auto-Negotiation

Q: Is 10/100Mbps auto-negotiation the same as Plug-n-Play?

No. The following configuration will cause severe network problems:



The hub, switch, or router will correctly sense (not auto-negotiate) the 10Mbps or 100Mbps speed. Since the end node was configured for a specific speed and duplex state and therefore does not negotiate, the hub, switch, or router will choose the communication mode specified by the 802.3u standard, namely half-duplex.

With one device running at half-duplex and the device on the other end of the connection at full-duplex, the connection will work reasonably well at low levels of traffic. At high levels of traffic the full-duplex device (end node, in this case) will experience an abnormally high level of CRC or alignment errors. The end users usually describe this situation as, "Performance seems to be approximately 1Mbps!" Often, end nodes will drop connections to their servers.

In this same situation, the half-duplex device will experience an abnormally high level of late collisions.

The network administrator must take care to verify the configuration of each network device during installation. Also, check the operational mode of each network device. That is, check both how you configured it and also that it comes up as you expect, for example, at 10Mbps/half-duplex.

Q: Is Gigabit Ethernet auto-negotiation the same as Plug-n-Play?

No. By the time the IEEE issued the 802.3z specification, they knew about the 10/100Mbps auto-negotiation problem (see the FAQ "Is 10/100Mbps auto-negotiation the same as Plug-n-Play?"). To prevent it, 802.3z auto-negotiation requires that, if one side of a connection is configured to auto-

negotiate, the other side must also auto-negotiate if the connection is to come up. In other words, if a switch is configured to auto-negotiate and its attached end node is configured to, say, 1000Mbps/full-duplex, the 803.2z spec requires that the switch NOT allow the link to come up.

Cabling

Q: How does the HP Auto-MDIX feature work? What products support this feature?

HP Auto-MDIX is a feature on all 10/100Base-TX ports of the HP ProCurve Networking Series 2300 and 2500 switches. When configured in the default configuration, "AUTO", the switch automatically detects the signaling on the cable from the connected device and operates as either an MDI or MDIX port. As a result, a "straight-through" twisted-pair cable can be used; you no longer have to use "crossover" cables, although "crossover" cables can also be used for any of the connections. The following devices support the HP Auto-MDIX feature:

- Fixed 10/100Base-T ports of the Switch 2512 and 2524 (J4812A and J4813A)
- Fixed 10/100Base-T ports of the Switch 2312 and 2324 (J4817A and J4818A)
- 100/1000-T Transceiver module (J4834A)

Q: Can I use category 3 cables with the HP ProCurve Switch 100/1000Base-T Transceiver?

No, only category 5 100-ohm UTP or STP cables are supported. In fact, for the most robust connections you should use cabling that complies with the Category 5E specifications, as described in Addendum 5 to the TIA-568-A standard (ANSI/TIA/EIA-568-A-5).

Q: What is the maximum length for cables used with the HP ProCurve Switch 100/1000Base-T Transceiver?

The maximum length is 100 meters using category 5, 100-ohm UTP or STP cable. This distance is correct for the IEEE 802.3ab specification.

General

Q: How can I set the time or date?

Use the global configuration mode **time** command in the Command Line Interface.

Features

Q: What are the differences between the HP ProCurve Switch 2512 (product number J4812A) and the HP ProCurve Switch 212M (product number J3298A)?

The HP ProCurve Switch 2512 has 12-fixed 10/100Base-T RJ-45 ports with two transceiver slots supporting Gigabit-SX, Gigabit-LX, 100/1000Base-T, 100-FX SC and Gigabit Stacking transceivers. The HP ProCurve Switch 212M has 12-fixed 10 Mbps RJ-45 ports with one transceiver slot supporting 10/100Base-T or 100Base-FX connectivity and one fixed 10/100Base-T RJ-45 port.

Q: What are the differences between the HP ProCurve Switch 2524 (product number J4813A) and the HP ProCurve Switch 2424M (product number J4093A)?

The HP ProCurve Switch 2524 has 24-fixed 10/100Base-T RJ-45 ports with two transceiver slots supporting Gigabit-SX, Gigabit-LX, 100/1000Base-T, 100-FX SC and Gigabit Stacking transceivers. The HP ProCurve Switch 2424M has 24-fixed 10/100Base-T RJ-45 ports with one module slot supporting Gigabit-SX, Gigabit-LX, 100/1000Base-T, 100-FX SC and Gigabit Stacking transceivers.

Q: What are the differences between the HP ProCurve Switch 2324 (product number J4818A) and the HP ProCurve Switch 2224 (product number J4095A)?

The HP ProCurve Switch 2324 is an unmanaged switch that has 24-fixed 10/100Base-T RJ-45 ports with two transceiver slots supporting Gigabit-SX, Gigabit-LX, 100/1000Base-T, 100-FX SC and Gigabit Stacking transceivers. The HP ProCurve Switch 2324 also provides a "download" port for future software upgrades. The HP ProCurve Switch 2224 is an unmanaged switch that has 24-fixed 10/100Base-T RJ-45 ports with one transceiver slot supporting 100Base-FX connectivity.

Q: What transceiver modules are supported on the HP ProCurve Networking Series 2300 and 2500 switches? Are the supported transceiver modules "hot-swappable?"

The following transceiver modules are supported for the Series 2300 and 2500 switches:

- HP ProCurve Networking Gigabit-SX Transceiver (J4131A)
- HP ProCurve Networking Gigabit-LX Transceiver (J4132A)
- HP ProCurve Networking 100-FX SC Transceiver (J4853A)
- HP ProCurve Networking 100/1000-T Transceiver (J4834A)

The transceivers that are part of the HP ProCurve Networking Gigabit Stacking Kit (J4116A)

Transceiver modules can be inserted into the switch while it is powered on. However, the switch must be rebooted to activate the transceiver. For more information, refer to chapter 2 of the HP ProCurve Networking Series 2300 and 2500 Switches *Installation and Getting Started Guide*.

Q: When I start a console session on my new HP ProCurve Networking Series 2500 switch, I get the prompt "HP ProCurve Switch 2524#." How do I access the menu interface?

The default interface for the HP ProCurve Networking Series 2500 switches is the Command Line Interface (CLI). The CLI is a full-featured interface that can be used to get status information (show commands) and perform configuration changes (configuration context). To access the menu interface, issue the command "menu" from the CLI. Also, using the "setup" command from the CLI or menu interface allows you to configure most of the basic options on the switch such as IP addressing and passwords.

Q: What function does the "download port" of the HP ProCurve Networking Series 2300 switch provide?

This port is used to connect a PC to the switch using a "null modem" cable. This connection can then be used to initiate a download of new software to the Series 2300 switch. New software may be available if there are any product enhancements, such as the support of new transceiver modules. For more information, refer to chapter 1 of the HP ProCurve Series 2300 and 2500 Switches *Installation and Getting Started Guide*.

Spanning Tree Protocol

Q: When STP is configured off, do the Switch 2512 and Switch 2524 forward STP bridge protocol data units (BPDUs) that they receive?

Yes. Note that the Switches 1600M, 2400M, 2424M, 4000M, and 8000M do not forward STP BPDUs when STP is configured off in those switches.

Q: Do the HP ProCurve Networking Series 2500 switches send 802.1Q tagged STP BPDUs out tagged ports?

No. These switch models implement single-instance STP and so their STP BPDUs are always untagged, regardless of the switch's configuration. For more information, see the HP ProCurve Networking Series 2500 Switches *Management and Configuration Guide* chapter on "Configuring Advanced Features", section "Effect of VLANs on Other Switch Features", subsection "Spanning Tree Protocol Operation with VLANs".

Q: What should I do if my switch (which has switched ports to the desktop) reports many Spanning Tree Protocol (STP) topology changes, even if I am not having connectivity problems?

STP was developed to manage switch-to-switch links, or bridge-to-bridge links as they were initially called. For most of STP's lifetime, a topology change was a serious event as it indicated a change in the LAN topology. It also indicated a likelihood that STP had to bring (at least) a portion of the LAN down for awhile in order to rebuild a new working topology. Of course, in a LAN we expect the topology to change only rarely.

Over the past few years, the cost of switch ports has dropped dramatically. Users can now afford to dedicate switch ports to end nodes.

Once an STP topology is stable, establishing an Ethernet link on a switch port results in a topology change. This happens every time a user powers up their PC (assuming that the PC is directly connected to a switch port), resets the PC, or brings up the PC's network stack. The establishment of this link causes the topology change count to increment in:

- The switch to which the end node is directly attached; and
- "Upstream" switches; that is, the root switch and the switches between the root switch and the switch to which the end node is directly connected.

This scenario does not result in any topology issues or changes. It does not result in any loss of connectivity in the LAN. If the incrementing of your topology change count is due to this type of scenario, you do not need to take any action.

The above discussion applies to Hewlett-Packard switches in general and HP ProCurve Networking switches in particular when configured in STP Normal mode. When an HP ProCurve Networking switch's port is configured in STP Fast mode, the switch will not increment the topology change count as a result of link changes on that port. Please see the switch's *Management and Configuration Guide* for details on normal and fast modes.

Q: When I power on my PC, I get the message "a file server could not be found." How do I fix that?

This is a well-known issue given the following situation:

- The PC is directly connected to a switch
- The PC is running Novell's VLMs or Client32
- The switch has Spanning Tree Protocol (STP) enabled

In this situation, when the directly connected PC is powered on, the switch senses linkbeat on that port. This causes the switch to go through the four Spanning Tree states: blocking, listening, learning, and forwarding. It takes 30 seconds for the switch to complete that sequence and begin forwarding packets to and from that port. During those 30 seconds, Novell sends 3 requests for a server, then stops looking. By the time Spanning Tree completes its job, Novell reports that "a file server could not be found."

There are several workarounds available:

- Disable Spanning Tree on the switch (if Spanning Tree is not needed, i.e. no loops in the network topology).

- For VLMs, add a "pause" just after calling VLM.EXE in STARTNET.BAT. When the user reboots a PC, have them wait at least 30 seconds before continuing the sequence. This workaround is documented on Novell's Knowledgebase (www.support.novell.com, search for document 2920460).
- For Client32, add a registry entry in the PC, as documented on Novell's Knowledgebase (search for document 2925582).

In the Fall of 1998, HP released switch firmware (version C.05.07 or greater) with an enhancement to resolve this timing problem between Novell and STP. The enhancement allows users to configure Spanning Tree so that it does not go through the 4 states, on a port-by-port basis. Instead, for those configured ports, Spanning Tree will immediately begin forwarding packets to and from the port. This allows Novell clients to communicate with the server as soon as the network card (NIC) is enabled. After that, the switch continues to listen for and send Spanning Tree packets on those configured ports. This protects the user who might inadvertently connect a hub or switch to that port and create a network loop—Spanning Tree will detect the loop after a short time, since the port listens for and sends STP packets on that port.

VLANs

Q: How can the Series 2500 switches allow an IP address per VLAN, but only one default gateway?

Visit the following HP ProCurve web page for the answer to this question:

http://www.hp.com/rnd/pdfs/switch_25xx_vlan_default_gateway_faq.pdf

Q: By default, VLAN support on the HP ProCurve Series 2500 switches is enabled. Can you disable VLAN support like you can on the HP ProCurve Switch 4000M?

No. VLAN support cannot be disabled on the HP ProCurve Series 2500 switches. By default, all ports are configured in the default VLAN (DEFAULT_VLAN). The following table shows the differences between the HP ProCurve Switch 4000M, 8000M, 1600M, 2424M and 2400M and the HP ProCurve Series 2500 switches with respect to VLAN support:

Feature	Switch 4000M, 8000M, 1600M, 2424M, and 2400M	Switch 2512 and 2524
VLAN Support Options	Disable (Default) Enable	Always enabled. (Cannot disable.)
Default VLAN	Can change VID	Cannot change VID.

Q: My HP ProCurve Networking Series 2500 switch is configured with multiple VLANs and I want to use DHCP to learn its IP address and default gateway for each VLAN. Since the HP ProCurve Networking Series 2500 switches support a single global default gateway, which of the DHCP supplied default gateway addresses will the switch use?

The Series 2500 switches will use the DHCP information received for the VLAN configured as the "Primary VLAN". In the factory-default configuration, the switch designates the default VLAN (DEFAULT_VLAN) as the "Primary VLAN". However, you can designate another VLAN as the "Primary VLAN" using the "Primary VLAN" configuration option. For more information, refer to chapter 9 of the HP ProCurve Series 2500 Switches *Management and Configuration Guide*.

Q: What is the recommended way to connect multiple VLANs between a routing switch and a layer 2 switch?

The diagram below illustrates the question.

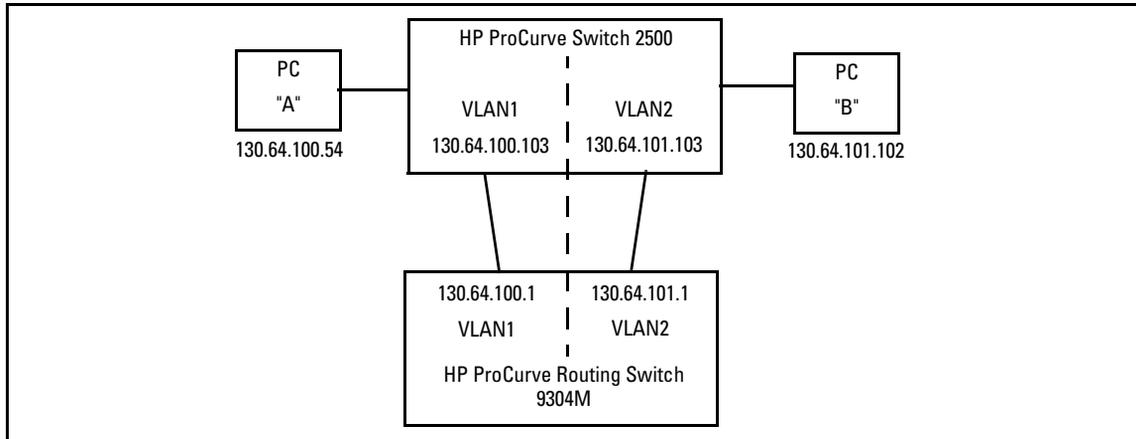


Figure 114. Example of Multiple VLANs

The following HP switches provide VLANs and have a single MAC/Ethernet address (filtering) table: Switch 800T, 2000, 1600M, 2400M, 2424M, 2512, 2524, 4000M, 8000M. In the diagram above we show a Switch 2500, but the following discussion applies to all of the switches listed in the previous sentence. The HP ProCurve Networking Routing Switch 9304M, 9308M, or 6308M-SX, as a default gateway, has a single MAC address (for all of its VLANs) if using virtual Ethernet interfaces. In the diagram above we show a 9304M, but this could be a 9308M or 6308M-SX as well.

Let's consider PC "A" attempting to send an IP packet to PC "B". PC "A" will send the 2500 a packet with the 9304M's MAC address in the destination field. If the 2500 has not yet learned this MAC address, the 2500 will flood the packet out all of its VLAN1 ports, including the VLAN1 link to the 9304M. The 9304M will then route the packet toward PC "B" via its link with the 2500's VLAN2 connection. The 2500 will enter the 9304M's MAC address into its MAC address table as located in VLAN2. The 2500 will also forward the packet to PC "B".

Let's consider a second packet that PC "A" sends to PC "B". PC "A" sends the packet, again addressed to the 9304M's MAC address, to the 2500. The 2500 will check its address table and find that the 9304M appears to be located on VLAN2. Since the 2500 believes that this MAC address is not located on VLAN1, the switch will discard the packet.

Later, when the 9304M transmits a packet to the 2500 via the VLAN1 link, the 2500 will update its address table to indicate that the 9304M's MAC address is located in VLAN1 instead of VLAN2. As you can see, the 2500's location information for the 9304M's MAC address will vary over time between VLAN1 and VLAN2. For this reason, some packets directed through the 2500 for the 9304M's MAC address will be discarded. Performance may appear to be poor or connectivity may appear to be broken.

To avoid this issue, simply use one cable between the 2500 and the 9304M instead of two, making sure that the two VLANs use tags on that link, as shown below.

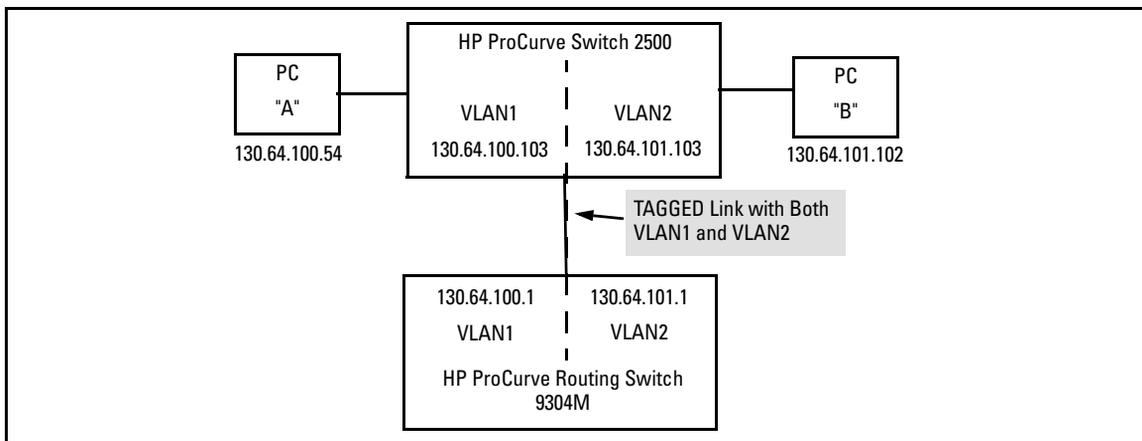


Figure 115. Example of Connecting Multiple VLANs Between a Routing Switch and a Layer 2 Switch

Troubleshooting

Q: Why won't my Xircom Realport-REM56G-100BTX Notebook NIC operate properly with my HP ProCurve Networking switch in full duplex mode at 100Mbps?

HP knows that this Xircom NIC, at full duplex over 100Mbps, does not operate with our HP ProCurve Networking switches. HP has no reports of any other vendors' NIC having a similar issue with our ProCurve switches. While this NIC's data sheet states support of full duplex, HP has not been able to get this NIC to operate at full duplex over a 100Mbps connection, regardless whether it is configured to Auto-negotiate or full duplex, fixed configuration.

HP is not aware of any fixes for this issue.

Failure symptoms include poor performance and lost connections. See the discussion of "Q: Is 10/100Mbps auto-negotiation the same as Plug-n-Play?" on this same page for more details on the problem symptoms.

The workaround is to operate the NIC in half-duplex mode or to use a different NIC.

Q: When I connected my new HP ProCurve Networking Series 2500 switch to my HP ProCurve Switch 4000M using a Gigabit-SX or Gigabit-LX connection, the link did not come up. Why not?

The HP ProCurve Networking Gigabit-SX module used in the HP ProCurve Switch 4000M, 8000M, 1600M, 2424M and 2400M is set by factory-default to "1000 Fdx" whereas the factory-default setting for the HP ProCurve Networking Gigabit-SX transceiver used in the HP ProCurve Networking Series 2300 and 2500 switches is "Auto". The configuration must be set to match on both ends to provide Gigabit connectivity.

Q: Using HP Hub & Switch Management for OV-UX (product number J3250N), how can I get the close-up view of my HP ProCurve Switch 2512 or 2524?

HP Hub & Switch Management for OV-UX, version N, displays these devices as generic icons. Double-clicking on these icons does not bring up the close-up view. You can use either of the following two workarounds.

- You can get the close-up view by entering the switch's IP address into the URL line (Internet Explorer uses the label "Address", while Netscape Communicator calls it "Location:") of a web browser. The supported browsers and versions are listed in the *Management and Configuration Guide* for the HP ProCurve Switch 2512 and 2524, in the chapter "Using the HP Web Browser Interface", section "Web Browser Interface Requirements".
- Configure Network Node Manager to bring up the switch's TELNET interface when double-clicking on the icon. To do this, perform the following steps:
 - a. Right-click on the switch icon.
 - b. From the menu, choose **Symbol Properties**.
 - c. Set Behavior to **Execute**.
 - d. Set **Application Action** to **Terminal Connect: Xterm TELNET**.
 - e. Click on **Target Objects**.
 - f. For **Additional Objects**, click **Add**.
 - g. Click **OK**.
 - h. Click **OK**.

Q: Why won't my iMac, which is directly attached to an HP ProCurve Networking switch, NetBoot from an OS X Server, even though it worked correctly when my iMac was plugged directly into a hub?

This issue is not specific to HP switches. Rather, it is a timing problem in early versions of the iMac Boot ROM. Apple Computer, Inc has fixed this NetBoot issue by releasing both an updated Boot ROM image and Mac OS X Server version 1.2. Please advise customers with iMac NetBoot issues to contact Apple Computer, Inc Technical Support.

Q: Why is my Macintosh system unable to use AppleTalk services?

Possible symptoms include: no AppleTalk services, only local network AppleTalk services, performance problems, and intermittent network services. If you remove the Macintosh from its dedicated switch (or routing switch or router) port and connect it to a hub, the problem goes away.

If the switch (or routing switch or router) has Spanning Tree Protocol enabled, see Apple Computer, Inc's Tech Info Library entry "Spanning Tree Protocol: AppleTalk Issues" at <http://til.info.apple.com/techinfo.nsf/artnum/n30922>



i n v e n t

The information contained in this document is subject to change without notice.

© 2001, 2003 Hewlett-Packard Company.
Reproduction, adaptation, or translation without prior written permission is prohibited except as allowed under the copyright laws.

HP Part Number: 5990-3102
Edition 2, June 2003