



Release Notes:

Version G.04.05 Operating System

for the HP Procurve Switch 4108GL

These release notes include information on the following:

- Downloading switch software and Documentation from the Web
- Enhancements in Release G.04.05
 - Friendly Port Names (page 6)
 - SSH Security—SSHv1 (page 11)
 - RADIUS Security (page 37)
 - Port-Access (802.1x) Security (page 65)
 - IP Preserve (page 91)
 - QoS Priority (page 94)
 - Terminating Remote Sessions (page 99)
 - Rapid Spanning-Tree—802.1W (page 121)
 - Port Security (page 101)
 - Fast-Uplink Spanning-Tree (page 133)
 - Show Tech Command (page 147)
- Software fix listings for the HP Procurve Switch 4108GL software releases (page 151)

Caution: Archive Pre-G.04.05 Configuration Files

A configuration file saved while using release G.04.05 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 G.04.05 before you update any switches to software release G.04.05 or later.

**© Copyright 2001-2002 Hewlett-Packard Company
All Rights Reserved.**

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

5990-3021
February 2002
Edition 2

Applicable Product

HP ProCurve Switch 4108GL (J4865A)

Trademark Credits

Microsoft, Windows, Windows 95, and Microsoft Windows NT are registered trademarks of Microsoft Corporation. Internet Explorer is a trademark of Microsoft Corporation. Ethernet is a registered trademark of Xerox Corporation. Cisco® is a trademark of Cisco Systems, Inc. Adobe® and Acrobat® are trademarks of Adobe Systems, Inc.

Software Credits

SSH in the HP Procurve Switch 4108GL 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 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 from the Web	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
Enhancements in Release G.04.05	5
Using Friendly (Optional) Port Names	6
Configuring and Operating Rules for Friendly Port Names	6
Configuring Friendly Port Names	7
Displaying Friendly Port Names with Other Port Data	8
Configuring Secure Shell (SSH)	11
Terminology	13
Prerequisite for Using SSH	13
Public Key Format Requirement	13
Steps for Configuring and Using SSH for Switch and Client Authentication	14
General Operating Rules and Notes	16
Configuring the Switch for SSH Operation	17
Further Information on SSH Client Public-Key Authentication	28
Messages Related to SSH Operation	33
Troubleshooting SSH Operation	34
Configuring RADIUS Authentication and Accounting	37
Terminology	38
Switch Operating Rules for RADIUS	38
General RADIUS Setup Procedure	39
Configuring the Switch for RADIUS Authentication	40
Configuring RADIUS Accounting	50
Operating Rules for RADIUS Accounting	51
Viewing RADIUS Statistics	56
Changing RADIUS-Server Access Order	61
Messages Related to RADIUS Operation	62
Troubleshooting RADIUS Operation	63

Configuring Port-Based Access Control (802.1x)	65
Why Use Port-Based Access Control (802.1x)?	65
General Features	65
How 802.1x Operates	66
Terminology	69
General Operating Rules and Notes	69
General Setup Procedure for Port-Based Access Control (802.1x)	70
Configuring Switch Ports as 802.1x Authenticators	72
Configuring Switch Ports To Operate As Supplicants for 802.1x Connections to Other Switches	78
Displaying 802.1x Configuration, Statistics, and Counters	81
How 802.1x Authentication Affects VLAN Operation	84
Messages Related to 802.1x Operation	87
Troubleshooting 802.1x Operation	88
IP Preserve: Retaining VLAN-1 IP Addressing Across Configuration File Downloads .	91
Operating Rules for IP Preserve	91
Configuring Port-Based Priority for Incoming Packets	94
Messages Related to Prioritization	98
Troubleshooting Prioritization	98
Using the "Kill" Command To Terminate Remote Sessions	99
Configuring and Monitoring Port Security	101
Basic Operation	101
Planning Port Security	103
CLI: Port Security Command Options and Operation	104
Web: Displaying and Configuring Port Security Features	113
Reading Intrusion Alerts and Resetting Alert Flags	113
Operating Notes for Port Security	119
Configuring Rapid Reconfiguration Spanning Tree (RSTP)	121
Overview	121
Transitioning from STP to RSTP	122
Configuring RSTP	123
Fast-Uplink Spanning Tree Protocol (STP)	133
Listing Switch Configuration and Operation Details for Help in Troubleshooting ...	147

Releases G.03.09, G.03.10, and G.03.13 149

Software Fixes 151

- Release G.03.09 (Beta Release Only) 151
- Release G.03.10 151
- Release G.03.13 152
- Release G.04.01 (Beta Release Only) 152
- Release G.04.02 (Beta Release Only) 153
- Release G.04.03 (Beta Release Only) 153
- Release G.04.04 (Beta Release Only) 153
- Release G.04.05 153

Index

Software Management

Caution: Archive Pre-G.04.05 Configuration Files

A configuration file saved while using release G.04.05 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 G.04.05 before you update any switches to software release G.04.05 or later.


Downloading Switch Documentation and Software from the Web

You can download software version G.04.05 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).
- HP's SNMP Download Manager included in HP TopTools for Hubs & Switches
- 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.

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 your switch.

TFTP Download from a Server

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

For example, to download an OS file named F_04_02.swi from a TFTP server with the IP address of 10.28.227.103:

1. Execute the copy command as shown below:

```
HP4108# copy tftp flash 10.28.227.103 F_04_02.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_02_03.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:

```
HP4108(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:

```
HP4108(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:

```
HP4108> 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 **write memory** from the Manager, Global, or Context configuration level.
- When exiting from the CLI to the Main Menu, press (for Yes) when you see the "save configuration" prompt:

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

Enhancements in Release G.04.05

Enhancement	Summary	Page
Friendly Port Names	Enables you to assign optional, meaningful names to physical ports on the switch.	6
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.	11
RADIUS	Protect access to the switch and monitor use of network resources through a centralized client authentication and accounting service.	37
Port-Access (802.1x)	Provide 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	65
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.)	91
QoS Priority	Enable assignment of non-default priority settings to inbound, untagged packets received on the switch.	94
Terminating Remote Sessions	Provides a kill command to terminate remote Telnet and SSH sessions.	99
Port Security	Provides port access control on the basis of device MAC addresses.	101
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.04.05 (or greater) release of your switch software	121
Fast-Uplink spanning tree (STP) mode for 802.1d spanning-tree operation	In a standard 802.1d spanning tree environment with redundant links, if the active link fails, the typical convergence time for a backup link to become the active, forwarding link is 30 seconds. Fast-uplink STP reduces the convergence time to approximately ten seconds.	133
show tech command	Outputs, in a single listing, switch operating and running configuration details from several internal switch sources.	147

Using Friendly (Optional) Port Names

Feature	Default	Menu	CLI	Web
Configure Friendly Port Names	Standard Port Numbering	n/a	page 7	n/a
Display Friendly Port Names	n/a	n/a	page 8	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 8.)
- 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 A3 on the switch to Bill Smith's workstation, and want to assign Bill's name and workstation IP address (10.25.101.73) as a port name for port A3:

```
HP4108(config)# int e A3 name Bill_Smith@10.25.101.73
HP4108(config)# write mem
HP4108(config)# show name A3
Port Names
  Port : A3
  Type : 10/100TX
  Name : Bill_Smith@10.25.101.73
```

Figure 1. Example of Configuring a Friendly Port Name

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

```
HP4108(config)# int e A5-A8 name Draft-Server:Trunk
HP4108(config)# write mem
HP4108(config)# show name 5-8
Port Names

  Port : A5
  Type : 10/100TX
  Name : Draft-Server:Trunk

  Port : A6
  Type : 10/100TX
  Name : Draft-Server:Trunk

  Port : A7
  Type : 10/100TX
  Name : Draft-Server:Trunk

  Port : A8
  Type : 10/100TX
  Name : Draft-Server:Trunk
```

Figure 2. 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:

```
HP4108(config)# show name
Port Names
Port Type      Name
-----
A1  10/100TX    not assigned
A2  10/100TX    not assigned
A3  10/100TX    Bill_Smith@10.25.101.73
A4  10/100TX    not assigned
A5  10/100TX    Draft-Server:Trunk
A6  10/100TX    Draft-Server:Trunk
A7  10/100TX    Draft-Server:Trunk
A8  10/100TX    Draft-Server:Trunk
A9  10/100TX    not assigned
A10 10/100TX    not assigned
A11 10/100TX    not assigned
A12 10/100TX    not assigned
.    .
.    .
.    .
```

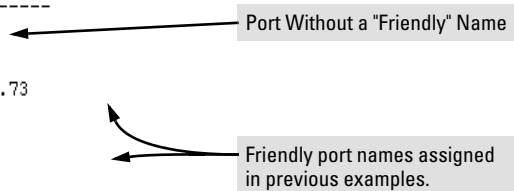


Figure 3. Example of Friendly Port Name Data for All Ports on the Switch

```
HP4108(config)# show name A2, A3, A5
Port Names
┌───┬───┬───┬───┬───┬───┐
| Port : A2 |
| Type : 10/100TX |
| Name : _not_assigned |
└───┴───┴───┴───┴───┴───┘
Port : A3
Type : 10/100TX
Name : Bill_Smith@10.25.101.73
Port : A5
Type : 10/100TX
Name : Draft-Server:Trunk
```

Figure 4. 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 A1 with the name "O'Connor_10.25.101.43", the show interface output for this port appears similar to the following:

```
HP4108(config)# show interface A1
Status and Counters - Port Counters for port A1

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
```

Figure 5. Example of a Friendly Port Name in a Per-Port Statistics Listing

Enhancements in Release G.04.05

Using Friendly (Optional) Port Names

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. (**show config** does not include ports that have only default settings in the startup-config file.)

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

For example, if you configure port A1 with a friendly port name:

```
HP4108(config)# int e A1 name Print_Server@10.25.101.43
HP4108(config)# write mem
HP4108(config)# int e A2 name Herbert's_PC
HP4108(config)# show config
Startup configuration:
; J4865A Configuration Editor; Created on release #G.04.XX
hostname "HP4108"
time daylight-time-rule None
no cdp run
interface A1
  name "Print_Server@10.25.101.43"
exit
snmp-server community "public" Unrestricted
vlan 1
  name "DEFAULT_VLAN"
  untagged 1-24
  ip address dhcp-bootp
  exit
no aaa port-access authenticator active
```

This command sequence saves the friendly port name for port A1 in the startup-config file, but does not do so for the name entered for port A2.

Listing includes friendly port name for port A1 only.

In this case, **show config** lists only port A1. Executing **write mem** after entering the name for port A2, and then executing **show config** again would result in a listing that includes both ports.

Figure 6. 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 18	n/a
Using the switch's public key	n/a	n/a	page 20	n/a
Enabling SSH	Disabled	n/a	page 22	n/a
Enabling client public-key authentication	Disabled	n/a	pages 25, 28	n/a
Enabling user authentication	Disabled	n/a	page 25	n/a

The Switch 4108GL uses 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. (See "PEM: (Privacy Enhanced Mode" on page 13.)

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.)

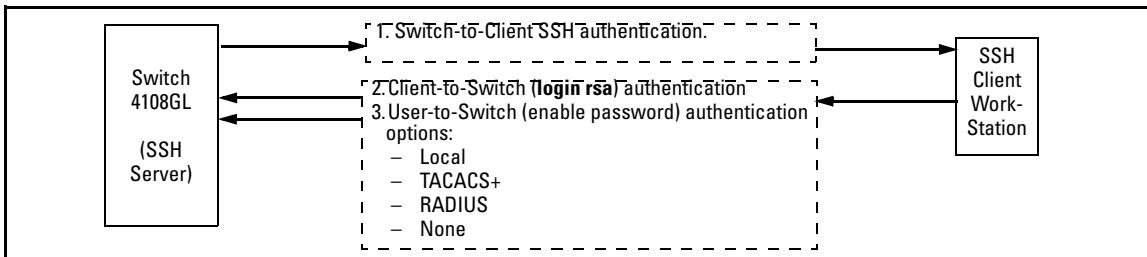


Figure 7. Client Public Key Authentication Model

Note

SSH in the HP Procurve Switch 4108GL 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 7. 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 7, 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.

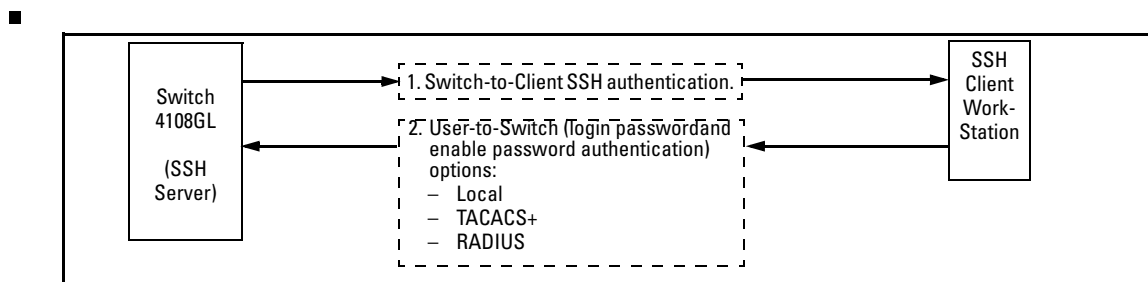


Figure 8. Switch/User Authentication

SSH on the Switch 4108GL 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 Switch 4108GL 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 9 and 10 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. A device's public key is used to authenticate the device to other devices.
- **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 18 and “4. Enabling SSH on the Switch and Anticipating SSH Client Contact Behavior” on page 22.)

Prerequisite for Using SSH

Before using a Switch 4108GL as an SSH server, you must first 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 11), 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 9. Example of Public Key in PEM-Encoded ASCII Format Common for SSHv2 Clients

```

512 37 78193303392019545793321845914508115859448079486918367079008218589443776362026267. . .
    
```

Figure 10. 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 1. 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 29.

B. Switch Preparation

1. Assign a login (Operator) and enable (Manager) password on the switch (page 18).
2. Generate a public/private key pair on the switch (page 18).

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 20).
4. Enable SSH on the switch (page 22).
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 28)

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 13.
- 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 24
show ip client-public-key [< babble fingerprint >]	page 31
show ip host-public-key [< babble fingerprint >]	page 21
show authentication	page 27
crypto key < generate zeroize > [rsa]	page 19
ip ssh	page 23
key-size < 512 768 1024 >	page 23
port < 1 - 65535 >	page 23
timeout < 5 .. 120 >	page 23
aaa authentication ssh	
login < local tacacs radius rsa >	page 25, 26
< local none >	page 25
enable < tacacs radius local >	page 25
< local none >	page 25
copy tftp pub-key-file <tftp server IP> <public key file>	page 31
clear public key	page 31

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 >

```
HP4108(config)# password all
New password for Operator: *****
Please retype new password for Operator: *****
New password for Manager: *****
Please retype new password for Manager: *****
HP4108(config)#
```

Figure 11. 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: crypto key generate [rsa]	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.)
crypto key zeroize [rsa]	Erases the switch's public/private key pair and disables SSH operation.
show ip ssh host-public-key [babble] [fingerprint]	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 21.) Displays a "fingerprint" of the switch's public key in hexadecimal format. (See "Displaying the Public Key" on page 21.)

For example, to generate and display a new key:

```
HP4108(config)# crypto key generate rsa
Generating new RSA host key.  If the cache is depleted,
this could take up to two minutes.
HP4108(config)# show ip host-public-key
896 35 4271994707660774263666250605799242148515279332487520218551264932934075407
04782860432930458032140273304999167004670769854352973485302001767770553555445568
80992231580238056056245444224389955500310200336191361046978602009243623264937429
4060627777506601747146563337525446401
```

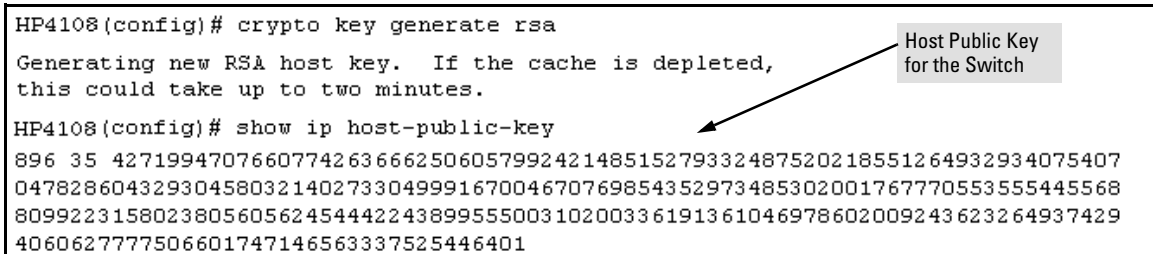


Figure 12. Example of Generating a Public/Private Host Key Pair for the Switch

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 13. 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 12.
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.

3. 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 14. Example of a Correctly Formatted Public Key (Unbroken ASCII String)

4. 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:

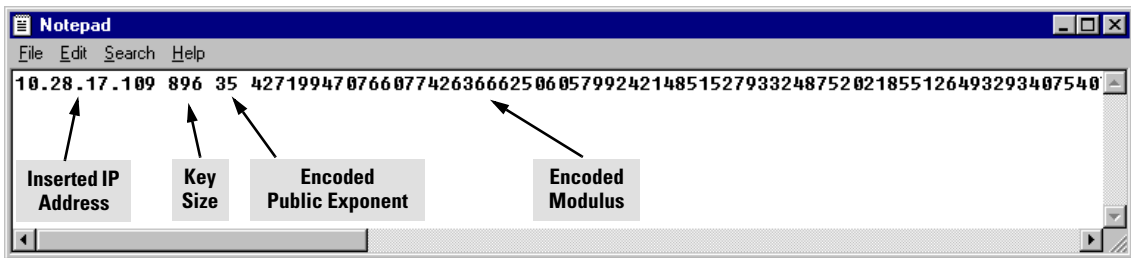


Figure 15. Example of a Switch Public Key Edited To Include the Switch's IP Address

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 14 on page 21.)
- **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 14 as follows:

```
HP4108# show ip host-public-key babble
896 xurac-metym-sagaf-recus-caleb-niten-kames-pobud-poluc-gelyl-exxsa

HP4108# show ip host-public-key fingerprint
896 f3:f3:61:4c:06:ea:53:c2:7f:e7:78:b6:9b:7f:88:5b
```

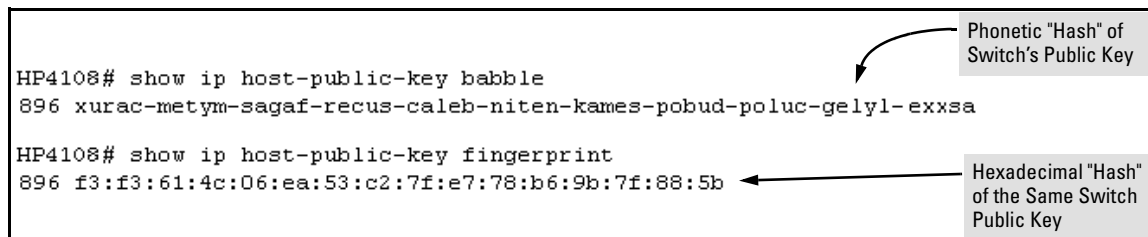


Figure 16. Examples of Visual Phonetic and Hexadecimal Conversions of the Switch's Public Key

Note

The two commands shown in figure 16 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 18.

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 25.

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 18.)
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 19).

Syntax: [no] ip ssh	Enables or disables SSH on the switch.
[key-size < 512 768 1024 >]	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.
[port < 1-65535 default >]	The IP port number for SSH connections (default: 22). Important: See the following "Note".
[timeout < 5 - 120 >]	The SSH login timeout value (default: 120 seconds).

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.

```
HP4108(config)# ip ssh
HP4108(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
```

The screenshot shows a terminal window with the following content:

- Command: `HP4108(config)# ip ssh` → Callout: "Enables SSH on the switch."
- Command: `HP4108(config)# show ip ssh` → Callout: "Lists the current SSH configuration and status."
- Output: `SSH Enabled : Yes`
- Output: `IP Port Number : 22` → Callout: "The switch uses these three settings internally for transactions with clients. See the **Note**, below."
- Output: `Timeout (sec) : 120`
- Output: `Server Key Size (bits) : 512`
- Table header: `Ses Type Source IP and Port`
- Table separator: `-----`
- Table rows:
 - `1 console`
 - `2 ssh 15.30.252.195:1722`
 - `3 telnet`
 - `4 inactive`

Callout for the table: "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 17. 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.

5. Configuring the Switch for SSH Authentication

Note that all methods in this section result in authentication of the switch's public key by an SSH client. However, only Option B, below results in the switch also authenticating the client's public key. Also, for a more detailed discussion of the topics in this section, refer to

Note

Hewlett-Packard recommends that you always assign a Manager-Level (enable) password to the switch. Without this level of protection, any user with Telnet, web, or serial port access to the switch can change the switch's configuration. *Also, if you configure only an Operator password, entering the Operator password through Telnet, web, or serial port access enables full manager privileges.* See "1. Assigning a Local Login (Operator) and Enable (Manager) Password" on page 18.

Option A: Configuring SSH Access for Password-Only SSH Authentication. When configured with this option, the switch uses its public key to authenticate itself to a client, but uses only passwords for client authentication.

Syntax: `aaa authentication ssh login < local | tacacs | radius >
[< local | none >]` Configures a password method for the primary and secondary login (Operator) access. If you do not specify an optional secondary method, it defaults to **none**.

`aaa authentication ssh enable < local | tacacs | radius >
[< local | none >]` Configures a password method for the primary and secondary enable (Manager) access. If you do not specify an optional secondary method, it defaults to **none**.

Option B: Configuring the Switch for Client Public-Key SSH Authentication. If configured with this option, the switch uses its public key to authenticate itself to a client, but the client must also provide a client public-key for the switch to authenticate. This option requires the additional step of copying a client public-key file from a TFTP server into the switch. This means that before you can use this option, you must:

1. Create a key pair on an SSH client.
2. Copy the client's public key into a public-key file (which can contain up to ten client public-keys).
3. Copy the public-key file into a TFTP server accessible to the switch and download the file to the switch.

(For more on these topics, refer to "Further Information on SSH Client Public-Key Authentication" on page 28.)

Enhancements in Release G.04.05

Configuring Secure Shell (SSH)

With steps 1-3, above, completed and SSH properly configured on the switch, if an SSH client contacts the switch, login authentication automatically occurs first, using the switch and client public-keys. After the client gains login access, the switch controls client access to the manager level by requiring the passwords configured earlier by the **aaa authentication ssh enable** command.

Syntax: copy tftp pub-key-file < ip-address > < filename >	Copies a public key file into the switch.
aaa authentication ssh login rsa < local none >	Configures the switch to authenticate a client public-key at the login level with an optional secondary password method (default: none).

Caution

To allow SSH access *only* to clients having the correct public key, you *must* configure the secondary (password) method for **login rsa** to **none**. Otherwise a client without the correct public key can still gain entry by submitting a correct local login password.

aaa authentication ssh enable < local tacacs radius > < local none >	Configures a password method for the primary and secondary enable (Manager) access. If you do not specify an optional secondary method, it defaults to none .
--	--

For example, assume that you have a client public-key file named Client-Keys.pub (on a TFTP server at 10.33.18.117) ready for downloading to the switch. For SSH access to the switch you want to allow only clients having a private key that matches a public key found in Client-Keys.pub. For Manager-level (enable) access for successful SSH clients you want to use TACACS+ for primary password authentication and **local** for secondary password authentication, with a Manager username of "leader" and a password of "m0ns00n". To set up this operation you would configure the switch in a manner similar to the following:


```

HP4108(config)# password manager user-name leader
New password for Manager: *****
Please retype new password for Manager: *****

HP4108(config)# aaa authentication ssh login rsa none
HP4108(config)# aaa authentication ssh enable tacacs local
HP4108(config)# copy tftp pub-key-file 10.33.18.117 Client-Keys.pub
HP4108(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 (enable) access. (Becomes available after SSH access is granted to a client.)

Figure 18. Configuring for SSH Access Requiring a Client Public-Key Match and Manager Passwords

Figure 19 shows how to check the results of the above commands.

```

HP4108(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	Local	None	Local	None
Port-Access	Local			
SSH	RSA	None	Tacacs	Local

```

HP4108(config)# show ip client-public-key
1024 35
114074066617014469079638036528401805391270437451114828825092785550110168603082616038
91468963065690359820412220255425432827643299433440329635043610210989476474605246515
64557222768203160764860366402053470340837100288429323150349226540935532111992745413
543765609589968291386053556814705585051025488575846923smith@fellow

512 35
92105526627259564161923578151309470522058178567194131272925848588434420808640403962
015631513328914504076264918047285352382655581475615051610646634228991kjlwilson@grayf
lds
  
```

Lists the current SSH authentication configuration.

Shows the contents of the public key file downloaded with the copy tftp command in figure 18. In this example, the file contains two client public-keys.

Figure 19. SSH Configuration and Client-Public-Key Listing From Figure 18

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 34 for possible solutions.

Further Information on SSH Client Public-Key Authentication

The section titled “5. Configuring the Switch for SSH Authentication” on page 25 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 25.)

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	11407406661701446907963803652840180539137043745111482882509285550110168603082603895914689630656903598204122202554254328276432994334403296350438102109894764746056455722276820316076486036640205347034083710028842932315034982365409355321119922465153140745413543765609589968291386053556814705585051025488575846923	smith@support.cairns.com

Figure 20. Example of a Client Public Key

Notes

Comments in public key files, such as `smith@support.cairns.com` in figure 20, 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 20 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 switch supports the following client-public-key properties:

Property	Supported Value	Comments
Key Format	ASCII (no PEM or other encoding)	See figure 14 on page 21. 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 20, 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:

```
HP4108(config)# copy tftp pub-key-file 10.38.252.195 clientkeys.txt
HP4108(config)# show ip client-public-key
1024 35
1140740666170144690796380365284018053912704374511148288250928555011016860308261603895
9146896306569035982041222025542543282764329943344032963504381021098947647460524651531
6455722276820316076486036640205347034083710028842932315034922654093553211199274541340
543765609589968291386053556814705585051025488575846923smith@fellow

512 35
9210552662725956416192357815130947052205817856719413127292584858843442080864040396207
015631513328914504076264918047285352382655581475615051610646634228991kjwilson@grayfile
lds
```

Figure 21. 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:

```
HP4108(config)# clear public-key
HP4108(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 24.</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 29 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>
Generating new RSA host key. If the cache is depleted, this could take up to two minutes.	<p>After you execute the crypto key generate [rsa] command, the switch displays this message while it is generating the key.</p>

Message	Meaning
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 33.

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 show ip host-public-key . If you see the message “XXXX”, 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 18.
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.
An attempt to copy a client public-key file into the switch has failed and the switch lists one of the following messages: Download failed: overlength key in key file. Download failed: too many keys in key file. Download failed: one or more keys is not a valid RSA public key.	The public key file you are trying to download has one of the following problems: A key in the file is too long. The maximum key length is 1024 characters, including spaces. This could also mean that two or more keys are merged together instead of being separated by a <CR><LF>. There are more than ten public keys in the key file. One or more keys in the file is corrupted or is not a valid rsa public key.

Symptom	Possible Cause
Client ceases to respond ("hangs") during connection phase.	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.

Configuring RADIUS Authentication and Accounting

Feature	Default	Menu	CLI	Web
Configuring RADIUS Authentication	None	n/a	page 40	n/a
Configuring RADIUS Accounting	None	n/a	page 50	n/a
Viewing RADIUS Statistics	n/a	n/a	page 56	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 Switch 4108GL:

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

Note

The switch does 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 49.

Accounting. RADIUS accounting on the switch 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 4108GL 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 4108GL, 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 56). 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 61.)
- 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 Switch 4108GL, 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 2. 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.

```
HP4108(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 22. 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 42
< console telnet ssh > < enable login > radius	page 42
< local none >	page 42
[no] radius-server host < <i>IP-address</i> >	page 44
[auth-port < <i>port-number</i> >]	page 44
[acct-port < <i>port-number</i> >]	page 44, 53
[key < <i>server-specific key-string</i> >]	page 44
[no] radius-server key < <i>global key-string</i> >	page 46
radius-server timeout < 1 .. 15 >	page 46
radius-server retransmit < 1 .. 5 >	page 46
[no] radius-server dead-time < 1 .. 1440 >	page 47
show radius	page 56
[< host < <i>ip-address</i> >]	page 57
show authentication	page 59
show radius authentication	page 59

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 50.)

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 11.

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):

```

HP4108(config)# aaa authentication telnet login radius none
HP4108(config)# aaa authentication telnet enable radius none
HP4108(config)# aaa authentication ssh login radius none
HP4108(config)# aaa authentication ssh enable radius none
HP4108(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 23. 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 48.

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 50 instead of continuing here.)

Syntax: [no] radius-server host < ip-address >	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 61.)
[auth-port < port-number >]	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)
[acct-port < port-number >]	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)
[key < key-string >]	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.
no radius-server host < ip-address > key	Use the no form of the command to remove the key for a specified server.

For example, suppose you have configured the switch as shown in figure 24 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".

```
HP4108(config)# show radius
Status and Counters - General RADIUS Information
  Deadtime(min) : 0
  Timeout(secs) : 5
  Retransmit Attempts : 5
  Global Encryption Key :
  Server IP Addr  Port  Port  Encryption Key
  -----
  10.33.18.127   1812 1813  TempKey01
```

Figure 24. Sample Configuration for RADIUS Server Before Changing the Key and Adding Another Server

To make the changes listed prior to figure 24, you would do the following:

```
HP4108(config)# radius-server host 10.33.18.127 key source0127
HP4108(config)# radius-server host 10.33.18.119 key source0119
HP4108(config)# show radius
Status and Counters - General RADIUS Information
  Deadtime(min) : 0
  Timeout(secs) : 5
  Retransmit Attempts : 5
  Global Encryption Key :
  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 24.

Figure 25. 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 61.

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 44.)
- **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 63.

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.

```
HP4108(config)# aaa authentication num-attempts 2
HP4108(config)# radius-server key My-Global-Key-1099
HP4108(config)# radius-server dead-time 5
HP4108(config)# radius-server timeout 3
HP4108(config)# radius-server retransmit 2
HP4108(config)# write mem
```

Figure 26. Example of Global Configuration Exercise for RADIUS Authentication

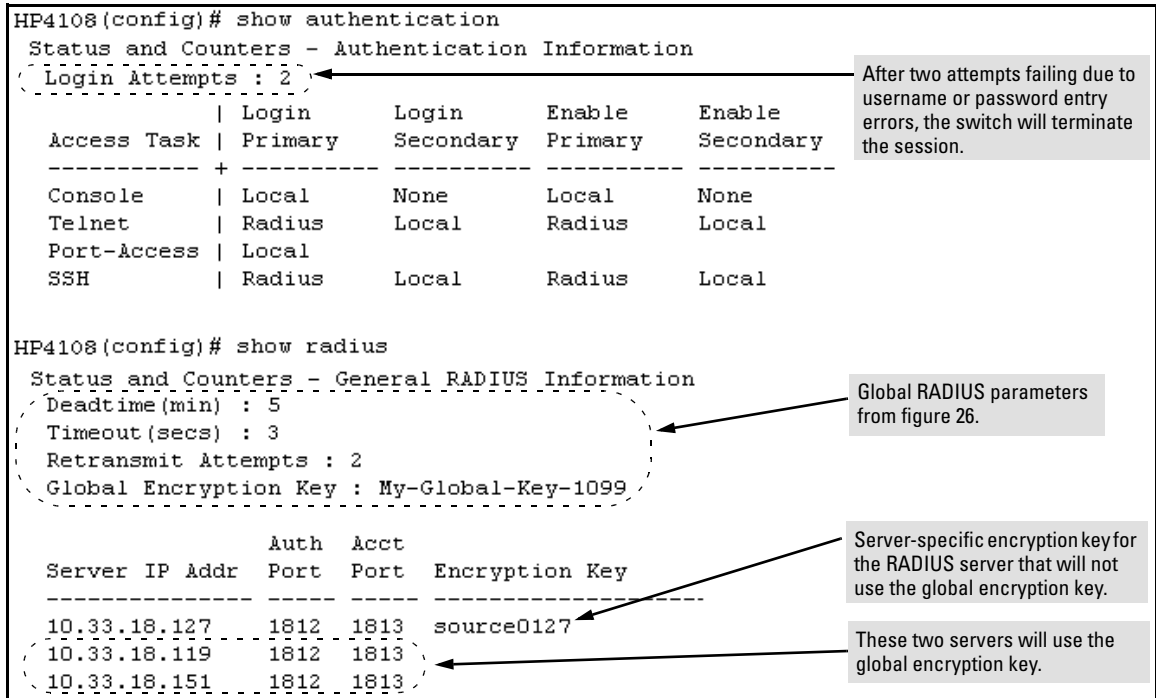


Figure 27. Listings of Global RADIUS Parameters Configured In Figure 26

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 53
[acct-port < port-number >]	page 53
[key < key-string >]	page 53
[no] aaa accounting < exec network system > < start-stop stop-only > radius	page 55
[no] aaa accounting update periodic < 1 .. 525600 > (in minutes)	page 56
[no] aaa accounting suppress null-username	page 56
show accounting	page 60
show accounting sessions	page 60
show radius accounting	page 60

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 39 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 Switch 4108GL supports 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 for the switch, 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 61.)
- 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 44.
- 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 44. (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 51), network (page 50), or system (page 51)
- **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 44. 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* > Adds a server to the RADIUS configuration or (with **no**) deletes a server from the configuration.

[acct-port < *port-number* >] 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)

[key < *key-string* >] 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.

(For a more complete description of the **radius-server** command and its options, turn to page 44.)

For example, suppose you want to 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.).

```
HP4108(config)# radius-server host 10.33.18.151 acct-port 1750 key source0151
HP4108(config)# write mem
HP4108(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
```

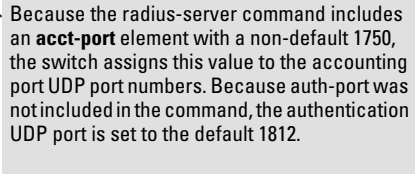


Figure 28. Example of Configuring for a RADIUS Server with a Non-Default Accounting UDP Port Number

The radius-server command as shown in figure 28, 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 37.)
- **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 37.) For information on this feature, refer to “Configuring Port-Based Access Control (802.1x)” on page 65.

Enhancements in Release G.04.05

Configuring RADIUS Authentication and Accounting

- **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 29, suppose that you wanted the switch to:

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

```
HP4108(config)# aaa accounting update periodic 10
HP4108(config)# aaa accounting suppress null-username

HP4108(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
```

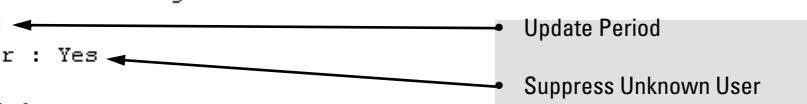


Figure 30. Example of Optional Accounting Update Period and Accounting Suppression on Unknown User

Viewing RADIUS Statistics

General RADIUS

Syntax: show radius 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.

[host < ip-addr >]

```

HP4108(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 31. Example of General RADIUS Information from Show Radius Command

```

HP4108(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 32. 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
show radius authentication

```
HP4108(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 33. Example of Authentication Information from the Show Authentication Command

```
HP4108 (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 34. Example of RADIUS Authentication Information from a Specific Server

RADIUS Accounting

Syntax: show accounting
 show radius accounting
 show accounting sessions

```
HP4108 (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 35. Example of the Accounting Configuration in the Switch

```
HP4108 (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 36. Example of RADIUS Accounting Information for a Specific Server

```
HP4108 (config)# show accounting sessions

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

Figure 37. 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:

```
HP4108(config)# show radius
Status and Counters - General RADIUS Information
  Deadtme(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 38. 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.

```

HP4108(config)# no radius host 10.10.10.003
HP4108(config)# no radius host 10.10.10.001
HP4108(config)# radius host 10.10.10.003
HP4108(config)# radius host 10.10.10.001

HP4108(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 39. 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

See also .

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>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>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>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)

  Server IP Addr  Auth Port  Acct 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

Configuring Port-Based Access Control (802.1x)

Feature	Default	Menu	CLI	Web
Configuring Switch Ports as 802.1x Authenticators	Disabled	n/a	page 72	n/a
Configuring Switch Ports to Operate as 802.1x Supplicants	Disabled	n/a	page 78	n/a
Displaying 802.1x Configuration, Statistics, and Counters	n/a	n/a	page 81	n/a
How 802.1x Affects VLAN Operation	n/a	n/a	page 84	n/a
RADIUS Authentication and Accounting	Refer to “Configuring RADIUS Authentication and Accounting” on page 37			

Why Use Port-Based Access Control (802.1x)?

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 Switch 4108GL 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.
 - 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.

Enhancements in Release G.04.05

Configuring Port-Based Access Control (802.1x)

- Use 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 the available authentication methods to only one: MD5.

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 40. Example of an 802.1x Application

Accounting . The Switch 4108GL also provide RADIUS Network accounting for 802.1x access. Refer to “Configuring RADIUS Accounting” on page 50.

How 802.1x Operates

Authenticator Operation

This operation provides security on a direct link between a single client and the switch, where both devices are 802.1x-aware. 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 A1 configured for 802.1x supplicant operation
- You want to connect port 1 on switch "A" to port B5 on switch "B".




Figure 41. Example of Supplicant Operation

1. When port A1 on switch "A" is first connected to a port on switch "B", or if the ports are already connected and either switch reboots, port A1 begins sending start packets to port B5 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 A1 receives a request packet from port B5, 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 A1 on switch "A".
3. Port A1 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 A1.
 - A "success" response unblocks port B5 to normal traffic from port A1.
 - A "failure" response continues the block on port B5 and causes port A1 to wait for the "held-time" period before trying again to achieve authentication through port B5.

Note

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

Terminology

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 Switch 4108GL 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 the Switch4108GL4 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, a 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.

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).

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.

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.
- 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.

Enhancements in Release G.04.05

Configuring Port-Based Access Control (802.1x)

- 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 active 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.
- You can configure a port as both an 802.1x authenticator *and* an 802.1x supplicant.
- 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.

Caution

To maintain security, you must disable LACP on all ports you intend to use for 802.1x port access. Otherwise, having both LACP and 802.1x port access enabled on a port creates a potential for a security breach.

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

Before You Begin

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.
3. 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.)

4. 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 of the steps for configuring 802.1x on the switch. For detailed information on each step, refer to “Configuring Switch Ports as 802.1x Authenticators” on page 72 or “Configuring Switch Ports To Operate As Supplicants for 802.1x Connections to Other Switches” on page 78.

1. Disable LACP on the ports on which you want to use 802.1x authentication. **Important: See the Caution on page 70.**
2. 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.
See page 72.
3. Configure the 802.1x authentication type. Options include:
 - Local Operator username and password (the default). This allows a client to use the switch's local username and password as valid 802.1x credentials for network access.
 - EAP RADIUS: Use if your RADIUS server application supports EAP authentication for 802.1x.
 - CHAP (MD5) RADIUS: Use if your RADIUS server application supports CHAP (MD5) authentication.See page 75.
4. If you selected 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 76.
5. Enable 802.1x authentication on the switch. See page 77.
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 (optional, step 7) on the switch, you should first ensure that the ports you have configured as 802.1x authenticators operate as expected.

Enhancements in Release G.04.05

Configuring Port-Based Access Control (802.1x)

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 76.
8. Configure 802.1x supplicant management on designated ports. See page 78.

Configuring Switch Ports as 802.1x Authenticators

802.1x Authentication Commands

[no] aaa port-access authenticator < [ethernet] < <i>port-list</i> >	page 72
[control quiet-period tx-period supplicant-timeout server-timeout max-requests reauth-period initialize reauthenticate clear-statistics]	page 72
aaa authentication port-access < local eap-radius chap-radius >	page 75
[no] aaa port-access authenticator active	page 77
[no] port-security [ethernet] < <i>port-list</i> > learn-mode port-access	page 76
802.1x Supplicant Commands	page 78
802.1x-Related Show Commands	page 81
RADIUS server configuration	page 37

1. Disable LACP on the Ports Selected for 802.1x Access

Syntax: no interface [e] < *port-list* > lacp Disables LACP on the designated ports.

Use **show lacp** to verify that LACP is disabled on the desired ports.

2. 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 77 to activate 802.1x authentication on the switch.)

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 77.

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

- [control < authorized | auto | unauthorized >] Controls authentication mode on a 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.
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.)
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.*
- [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 the server does not respond 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)

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

[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)

[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)

[initialize]

On the specified ports, blocks inbound and outbound traffic and restarts the 802.1x authentication process. This behavior occurs 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.

2. 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
 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 application.)
- chap-radius
 Use CHAP-RADIUS 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:

```

HP4108(config)# aaa authentication port-access eap-radius
HP4108(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 42. Example of 802.1x (Port-Access) Authentication

3. 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 37.

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.

4. Optional: 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 action < none send-alarm send-disable >	Configures port-security on the specified port(s) to allow only the first 802.1x-aware device that the port detects.
--	--

Note

Port-Security operates with 802.1x authentication as described above only if the affected 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 A10 for 802.1x authenticator operation and display the result:

```
HP4108(config)# aaa port-access authenticator e A10 control auto
HP4108(config)# show port-access authenticator e A10 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:

<pre>aaa port-access authenticator e < port-list ></pre>	Enables 802.1x authentication on the port.
<pre>aaa port-access authenticator e < port-list > control auto</pre>	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:

<pre>aaa port-access authenticator e < port-list > control auto</pre>	Forces the port to accept only a device that supports 802.1x and supplies valid credentials.
---	--

5. Enable 802.1x Authentication on the Switch

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

Syntax:

```
aaa port-access authenticator active
```

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

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

802.1x Authentication Commands	page 72
802.1x Supplicant Commands	
[no] aaa port-access < supplicant < [ethernet] < <i>port-list</i> >	page 79
[auth-timeout held-period start-period max-start initialize identity secret clear-statistics]	page 80
802.1x-Related Show Commands	page 81
RADIUS server configuration	pages 37

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 A1 configured for 802.1x supplicant operation
- You want to connect port A1 on switch "A" to port B5 on switch "B".




Figure 43. Example of Supplicant Operation

1. When port A1 on switch "A" is first connected to a port on switch "B", or if the ports are already connected and either switch reboots, port A1 begins sending start packets to port B5 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 A1 receives a request packet from port B5, 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 75), 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 A1 on switch "A".
 3. Port A1 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 A1.
 - A "success" response unblocks port B5 to normal traffic from port A1.
 - A "failure" response continues the block on port B5 and causes port A1 to wait for the "held-time" period before trying again to achieve authentication through port B5.

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* >]

[**secret**]

Enter secret: < *password* >

Repeat secret: < *password* >

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)

[**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 78 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).

Syntax (Continued from page 80):

[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 <i>start-period</i> for a response. If no response comes during the <i>start-period</i> , 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 72
802.1x Supplicant Commands	page 78
802.1x-Related Show Commands	
show port-access authenticator	below
show port-access supplicant	page 83
RADIUS server configuration	pages 37

Show Commands for Port-Access Authenticator

Syntax: show port-access authenticator	Shows 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.
[e] < port-list >	Same as above, but limits port status to only the specified port. The statistics values are blank if the specified port is not enabled as an authenticator.

Syntax (Continued from page 81):

config	Shows whether port-access authenticator is active and the 802.1x configuration of the specified port. The configuration settings are blank if the specified port is not enabled as an authenticator.
statistics	Shows whether port-access authenticator is active and the statistics of the specified port. Includes the supplicant's MAC address, as determined by the content of the last EAPOL frame received on the port. The statistics values are blank if the specified port is not enabled as an authenticator.
show port-access authenticator	
[e] < port-list >	
session-counters	Shows whether port-access authenticator is active the session data, session status on the specified port. 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 80.) The fields are blank if the specified port is not enabled as an authenticator.
[config]	Same as the [e] < port-list > config command (above), but for all ports on the switch that are enabled as authenticators.
[e] < port-list >	Same as the [e] < port-list > config command (above).
[statistics]	Same as the statistics command (above), but for all ports on the switch that are enabled as authenticators.
[e] < port-list >	Same as the [e] < port-list > statistics command (above).
[session-counters]	Same as the [e] < port-list > session-counters command (above), but for all ports on the switch that are enabled as authenticators.
[e] < port-list >	Same as the [e] < port-list > session-counters command (above).

Show Commands for Port-Access Supplicant

show port-access supplicant	Shows the port-access supplicant configuration (excluding the secret parameter) for the 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 80). For descriptions of the supplicant parameters, refer to "Configuring a Supplicant Switch Port" on page 80.
[e] < port-list >	Same as the above command, but for the specified port(s). If a port is not configured as a supplicant, it does not appear in the listing.
[statistics]	Shows the port-access statistics and source MAC address(es) for all ports configured on the switch as supplicants. See the "Note", below.
[e] < port-list >	Same as the above statistics command, but for the specified port(s). If a port is not configured as a supplicant, it does not appear in the listing.

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 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 most recent transaction statistics 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 802.1x Authentication Affects VLAN Operation

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.)

Static VLAN Requirement

The static VLAN to which a client is assigned must already exist on the switch. If it does not exist or is a dynamic VLAN (created by GVRP), authentication will fail. 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 a Port Is Not 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 that the RADIUS server specifies, 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 A2 requires access to VLAN 22, but VLAN 22 is configured for no access on port A2, and VLAN 33 is configured as untagged on port A2:

```
HP4108
----- CONSOLE - MANAGER MODE -----
Switch Configuration - VLAN - VLAN Port Assignment

Port  default_vlan  vlan_22  vlan_33  vlan_44
----+-----
A1   | Untagged   Tagged   No       No
A2   | No         No       Untagged No
A3   | Untagged   Forbid   Forbid   Forbid
A4   | 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 action.
```

Scenario: An authorized 802.1x client requires access to VLAN 22 from port A2. However, access to VLAN 22 is blocked (not untagged or tagged) on port A2 and VLAN 33 is untagged on port A2.

Figure 44. Example of an Active VLAN Configuration

In figure 44, 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.
- 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.

```
HP4108(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
-----
A1              Tagged      Learn          Up
A2              802.1X      Learn          Up
A4              Tagged      Learn          Up
.              .          .              .
.              .          .              .
.              .          .              .

Port VLAN membership configuration

Port  Mode
----  ---
A1    Tagged
A2    No
A4    Tagged
```

This entry shows that port A2 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 44), the only time port A2 appears in this **show vlan 22** listing is during an 802.1x session with an attached client. Otherwise, port A2 is not listed.

Figure 45. The Active Configuration for VLAN 22 Temporarily Changes for the 802.1x Session

Enhancements in Release G.04.05
 Configuring Port-Based Access Control (802.1x)

- With the preceding in mind, since (static) VLAN 33 is configured as untagged on port A2 (see figure 44), and since a port can be untagged on only one VLAN, port A2 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.

Even though port A2 is configured as Untagged on (static) VLAN 33 (see figure 44), it does not appear in the VLAN 33 listing while the 802.1x session is using VLAN 22 in the Untagged status. However, after the 802.1x session with VLAN 22 ends, the active configuration returns port A2 to VLAN 33.

```

HP4108# show vlan 33

Status and Counters - VLAN Information - Ports - VLAN 33

802.1Q VLAN ID : 33
Name           : VLAN_33
Status        : Static

Port Information Mode      Unknown VLAN Status
-----
A4              Tagged    Learn          Up

Port VLAN membership configuration

Port Mode
-----
A2  Untagged
    
```

Figure 46. The Active Configuration for VLAN 33 Temporarily Drops Port 22 for the 802.1x Session

When the 802.1x client's session on port A2 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 A2 ends, VLAN 22 access on port A2 also ends, and the untagged VLAN 33 access on port A2 is restored.

After the 802.1x session on VLAN 22 ends, the active configuration again includes VLAN 33 on port 2.

```

HP4108# show vlan 33

Status and Counters - VLAN Information - Ports - VLAN 33

802.1Q VLAN ID : 33
Name           : VLAN_33
Status        : Static

Port Information Mode      Unknown VLAN Status
-----
A2              Untagged  Learn          Down
A4              Tagged    Learn          Down
    
```

Figure 47. The Active Configuration for VLAN 33 Restores Port A2 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

Message	Meaning
Port < <i>port-list</i> > 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: HP4108(config)# aaa port-access authenticator e 10
Port < <i>port-list</i> > 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 79.
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 < <i>x.x.x.x</i> >, try the suggestions listed for that message (page 62).

Troubleshooting 802.1x Operation

Note

To list the 802.1x port-access Event Log messages stored on the switch, use **show log 802**.

See also “Troubleshooting RADIUS Operation” on page 63.

Symptom	Possible Cause
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).	There can be several reasons for not receiving a response to an authentication request. Do the following: <ul style="list-style-type: none">• Use ping to ensure that the switch has access to the configured RADIUS servers.• Verify that the switch is using the correct encryption key (RADIUS secret key) for each server.• Verify that the switch has the correct IP address for each RADIUS server.• Ensure that the radius-server timeout period is long enough for network conditions.
The switch does not authenticate a client even though the RADIUS server is properly configured and providing a response to the authentication request.	If the RADIUS server configuration for authenticating the client includes a VLAN assignment, ensure that the VLAN exists as a static VLAN on the switch. Refer to “How 802.1x Authentication Affects VLAN Operation” on page 84.
During RADIUS-authenticated client sessions, access to a VLANs on the port used for the client sessions is lost.	If the affected VLAN is configured as untagged on the port, it may be temporarily blocked on that port during an 802.1x session. This is because the switch has temporarily assigned another VLAN as untagged on the port to support the client access, as specified in the response from the RADIUS server. Refer to “How 802.1x Authentication Affects VLAN Operation” on page 84.
The switch appears to be properly configured as a supplicant, but cannot gain access to the intended authenticator port on the switch to which it is connected.	If aaa authentication port-access is configured for Local, ensure that you have entered the local <i>login</i> (operator-level) username and password of the authenticator switch into the identity and secret parameters of the supplicant configuration. If instead, you enter the enable (manager-level) username and password, access will be denied.
The supplicant statistics listing shows multiple ports with the same authenticator MAC address.	The link to the authenticator may have been moved from one port to another without the supplicant statistics having been cleared from the first port. Refer to “Note on Supplicant Statistics” on page 83.
The show port-access authenticator <port-list> command shows one or more ports remain open after they have been configured with control unauthorized .	802.1x is not active on the switch. After you execute aaa port-access authenticator active , all ports configured with control unauthorized should be listed as Closed.

Symptom	Possible Cause
---------	----------------

```

HP4108(config)# show port-access authenticator e A9
Port Access Authenticator Status
  Port-access authenticator activated [No] : No
          Access  Authenticator  Authenticator
Port Status Control  State          Backend State
-----
A9  Open  FU          Force Auth  Idle

HP4108(config)# aaa port-access authenticator active

HP4108(config)# show port-access authenticator e A9
Port Access Authenticator Status
  Port-access authenticator activated [No] : Yes
          Access  Authenticator  Authenticator
Port Status Control  State          Backend State
-----
A9  Closed FU          Force Unauth Idle
  
```

Port A9 shows an "Open" status even though Access Control is set to **Unauthorized (Force Auth)**. This is because the port-access authenticator has not yet been activated.

RADIUS server fails to respond to a request for service, even though the server's IP address is correctly configured in the switch.

Use **show radius** to verify that the encryption key (RADIUS secret key) the switch is using is correct for the server being contacted. If the switch has only a global key configured, 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.

```

10.33.18.119(config)# show radius
Status and Counters - General RADIUS Information
  Deadtme(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

Also, ensure that the switch port used to access the RADIUS server is not blocked by an 802.1x configuration on that port. For example, **show port-access authenticator <port-list>** gives you the status for the specified ports. Also, ensure that other factors, such as port security or any 802.1x configuration on the RADIUS server are not blocking the link.

Enhancements in Release G.04.05
Configuring Port-Based Access Control (802.1x)

Symptom	Possible Cause
<p>The authorized MAC address on a port that is configured for both 802.1x and port security either changes or is re-acquired after execution of aaa port-access authenticator <port-list> initialize.</p>	<p>If the port is force-authorized with aaa port-access authenticator <port-list> control authorized command and port security is enabled on the port, then executing initialize causes the port to clear the learned address and learn a new address from the first packet it receives after you execute initialize.</p>
<p>A trunked port configured for 802.1x is blocked.</p>	<p>If you are using RADIUS authentication and the RADIUS server specifies a VLAN for the port, the switch allows authentication, but blocks the port. To eliminate this problem, either remove the port from the trunk or reconfigure the RADIUS server to avoid specifying a VLAN.</p>

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

IP Preserve enables you to copy a configuration file to multiple Switch 4108GLs 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 48, 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).

```

; J4865A Configuration Editor; Created on release #G.04.01
hostname "HP4108"
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 48. Example of Implementing IP Preserve in a Configuration File

Enhancements in Release G.04.05

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

For example, consider Figure 49:

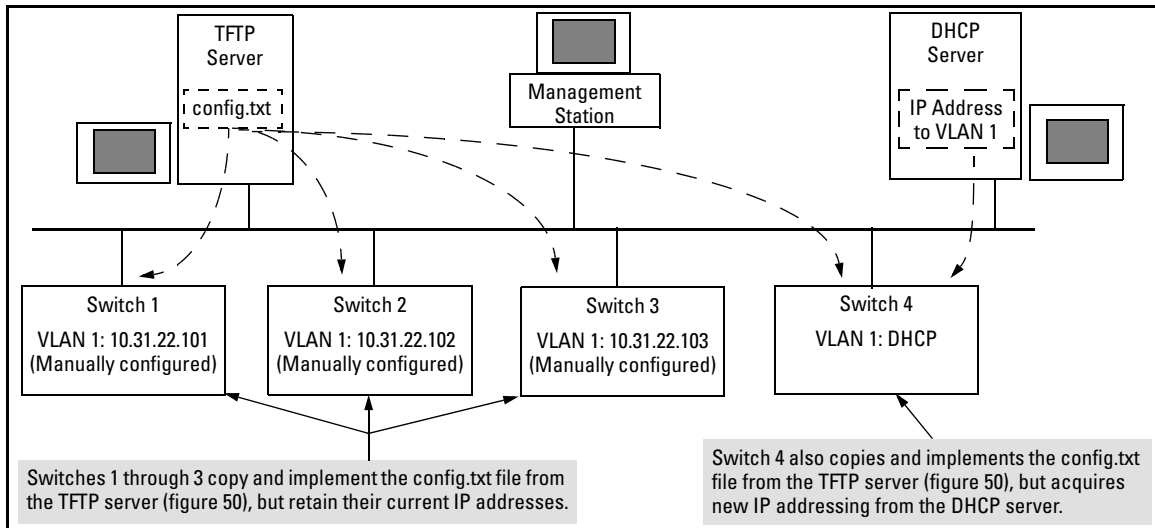


Figure 49. Example of IP Preserve Operation

If you apply the following configuration file to Figure 49, 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.

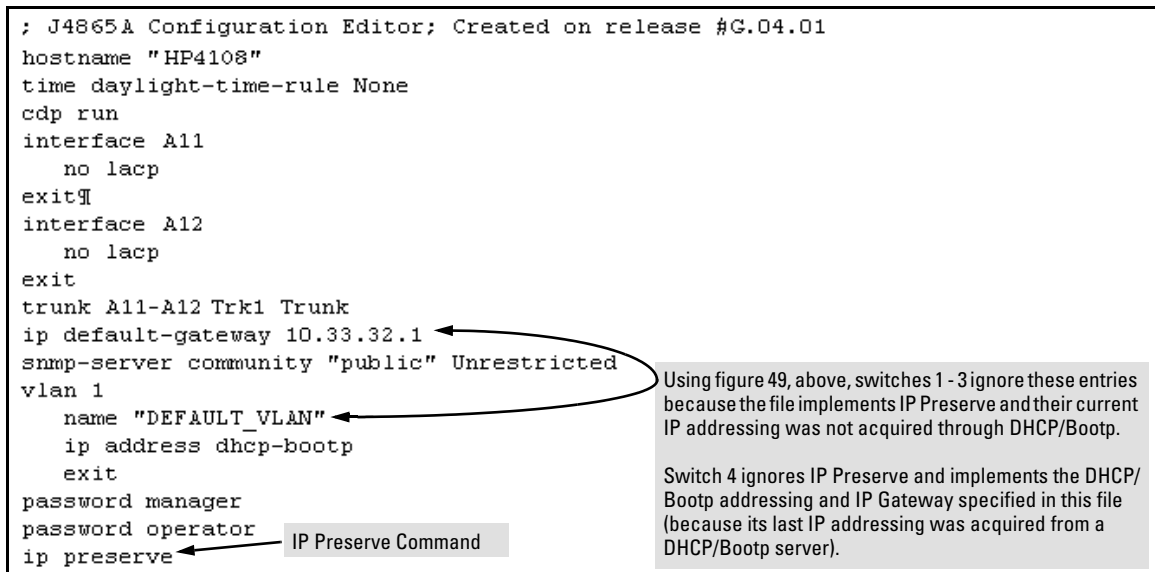


Figure 50. Configuration File in TFTP Server, with DHCP/Bootp Specified as the IP Addressing Source

If you apply this configuration file to figure 49, 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.

```

; J4865A Configuration Editor; Created on release #G.04.01

hostname "HP4108"
time daylight-time-rule None
cdp run
interface A11
  no lACP
exit
interface A12
  no lACP
exit
trunk A11-A12 Trk1 Trunk
ip default-gateway 10.33.32.1
snmp-server community "public" Unrestricted
vlan 1
  name "DEFAULT_VLAN"
  forbid A3
  untagged A1, A7-A10, A13-A14, Trk1
  tagged A4-A6
  no untagged A2-A3
  ip address 10.31.22.255 255.255.248.0
  exit
password manager
password operator
ip preserve

```

Because switch 4 (figure 49) 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 51. 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 96	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
- Inbound, tagged VLAN packets having a priority level of 0 (zero)

(The switch does not alter the existing priority level of inbound, tagged VLAN packets carrying a priority level of 1-7.)

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 traffic and 0-priority tagged 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 priority setting 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

Switch 4108GL ports use three outbound port queues, *Low*, *Normal*, and *High*. As described below, these three queues map to the eight priority settings specified in the 802.1p standard.

Table 3. Mapping Priority Settings to Device Queues

802.1p Priority Settings Used In Tagged VLAN Packets	Switch 4108GL Outbound Port Queues	Queue Assignment in Downstream Devices With:			
		8 Queues	4 Queues	3 Queues	2 Queues
1 (low)	Low	1	1	1	1
2 (low)	Low	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

For example, suppose you have configured port A10 to assign a priority level of 1 (low):

- An untagged packet coming into the switch on port A10 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 an 802.1p priority setting of 0 (zero) coming into the switch on port A10 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 an 802.1p priority setting (1 - 7) coming into the switch on port A10 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 A10).

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 the Switch 4108GL

- 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 - 7, an inbound, *untagged* packet adopts the specified priority and is sent to the corresponding outbound queue on the outbound port. (See table 3, "Mapping Priority Settings to Device Queues", on page 95.) 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, when port-based priority is configured as 0 - 7, an inbound, *tagged* packet with a priority of 0 (zero) adopts the specified priority and is sent to the corresponding outbound queue on the outbound port. (See table 3, "Mapping Priority Settings to Device Queues", on page 95.) 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 with a preset priority of 0 - 7 in its tag keeps that priority. It is assigned an outbound queue on the basis of that priority (regardless of the port-based priority configured on the port). (Refer to table 3, "Mapping Priority Settings to Device Queues" on page 95.)
- 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 >	Configures a non-default port-based 802.1p priority for incoming, untagged packets or tagged packets arriving with a "0" priority on the designated ports, as described under "Operating Rules for Port-Based Priority", above.
qos priority 0	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 3 on page 95.)
show running-config	Lists any non-default (1 - 7) port-based priority settings in the running-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.

show config

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.

For example, suppose you wanted to configure ports A10-A12 on the switch to prioritize all untagged, inbound VLAN traffic as "Low" (priority level = 1; refer to table 3 on page 95).

```
HP4108(config)# interface e A9-A12 qos priority 1
HP4108(config)# write mem
HP4108(config)# show config
Startup configuration:
; J4865A Configuration Editor; Created on release #G.04.04
hostname "HP4108"
time daylight-time-rule None
_cdp_run_ _ _ _ _
| interface A9          |
|   qos priority 1     |
| exit                 |
| interface A10         |
|   qos priority 1     |
| exit                 |
| interface A11         |
|   qos priority 1     |
| exit                 |
| interface A12         |
|   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 A9 -A12 to "1" (Low) and saves the configuration changes to the startup-config file.

Ports A9 - A12 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 52. Example of Configuring Non-Default Prioritization on Untagged, Inbound Traffic

Messages Related to Prioritization

Message	Meaning
<code>< priority-level >: Unable to create.</code>	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:

```
HP4108(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

HP4108(config)# kill 2
HP4108(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
```

Session 2 is an active Telnet session.

The kill 2 command terminates session 2.

Figure 53. Example of Using the "Kill" Command To Terminate a Remote Session

Enhancements in Release G.04.05
Using the "Kill" Command To Terminate Remote Sessions

Configuring and Monitoring Port Security

Feature	Default	Menu	CLI	Web
Displaying Current Port Security	n/a	—	page 107	page 113
Configuring Port Security	disabled	—	page 108	page 113
Intrusion Alerts and Alert Flags	n/a	page 118	page 116	page 119

Using Port Security, you can configure each switch port with a unique list of the MAC addresses of devices that are authorized to access the network through that port. This enables individual ports to detect, prevent, and log attempts by unauthorized devices to communicate through the switch.

Note

This feature does not prevent intruders from receiving broadcast and multicast traffic.

Basic Operation

Default Port Security Operation. The default port security setting for each port is off, or “continuous”. That is, any device can access a port without causing a security reaction.

Intruder Protection. A port that detects an “intruder” blocks the intruding device from transmitting to the network through that port.

General Operation for Port Security. On a per-port basis, you can configure security measures to block unauthorized devices, and to send notice of security violations. Once you have configured port security, you can then monitor the network for security violations through one or more of the following:

- Alert flags that are captured by network management tools such as HP TopTools for Hubs & Switches
- Alert Log entries in the switch’s web browser interface
- Event Log entries in the console interface
- Intrusion Log entries in either the menu interface, CLI, or web browser interface

For any port, you can configure the following:

- **Authorized (MAC) Addresses:** Specify up to eight devices (MAC addresses) that are allowed to send inbound traffic through the port. This feature:
 - Closes the port to inbound traffic from any unauthorized devices that are connected to the port.
 - Provides the option for sending an SNMP trap notifying of an attempted security violation to a network management station and, optionally, disables the port. (For more on configuring the switch for SNMP management, see “Trap Receivers and Authentication Traps” in the *Management and Configuration Guide* for your switch.)

Blocking Unauthorized Traffic

Unless you configure the switch to disable a port on which a security violation is detected, the switch security measures block unauthorized traffic without disabling the port. This implementation enables you to apply the security configuration to ports on which hubs, switches, or other devices are connected, and to maintain security while also maintaining network access to authorized users. For example:

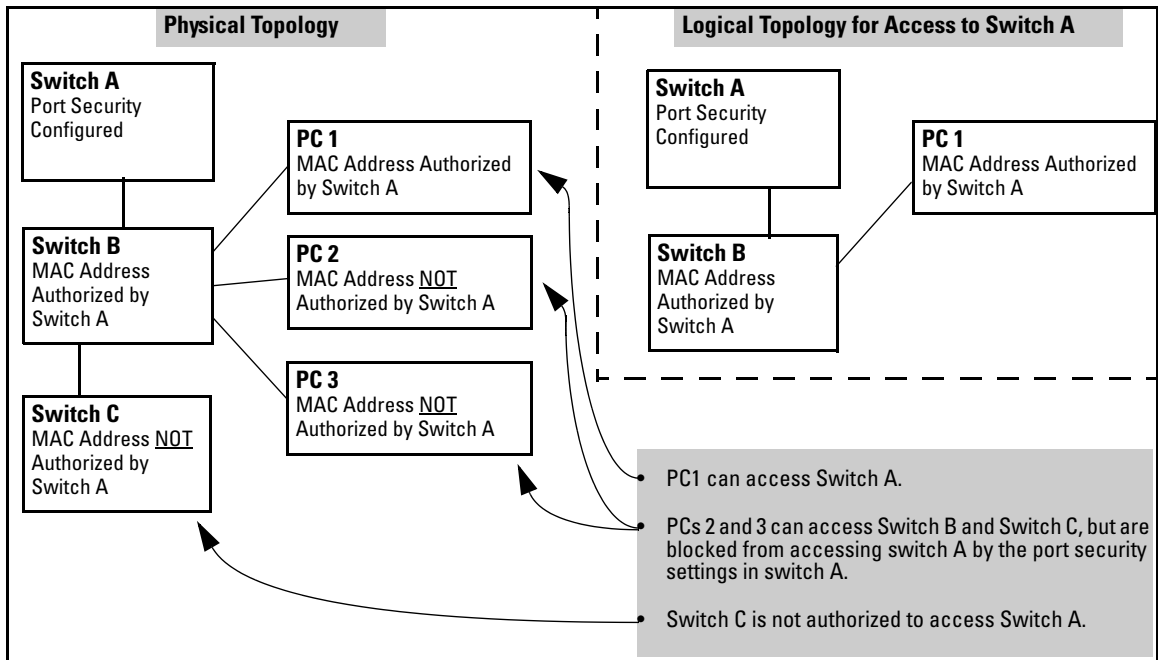


Figure 54. Example of How Port Security Controls Access

Note

Broadcast and Multicast traffic is not “unauthorized” traffic, and can be read by intruders connected to a port on which you have configured port security.

Trunk Group Exclusion

Port security does not operate on either a static or dynamic trunk group. If you configure port security on one or more ports that are later added to a trunk group, the switch will reset the port security parameters for those ports to the factory-default configuration. (Ports configured for either Active or Passive LACP, and which are not members of a trunk, can be configured for port security.)

Planning Port Security

1. Plan your port security configuration and monitoring according to the following:
 - a. On which ports do you want to configure port security?
 - b. Which devices (MAC addresses) are authorized on each port (up to 8 per port)?
 - c. For each port, what security actions do you want? (The switch automatically blocks intruders detected on that port from transmitting to the network.) You can configure the switch to (1) send intrusion alarms to an SNMP management station and to (2) optionally disable the port on which the intrusion was detected.
 - d. How do you want to learn of the security violation attempts the switch detects? You can use one or more of these methods:
 - Through network management (That is, do you want an SNMP trap sent to a net management station when a port detects a security violation attempt?)
 - Through the switch’s Intrusion Log, available through the CLI, menu, and web browser interface
 - Through the Event Log (in the menu interface or through the CLI **show log** command)
2. Use the CLI or web browser interface to configure port security operating and address controls. The following table describes the parameters.

CLI: Port Security Command Options and Operation

Port Security Commands Used in This Section

show port-security	page 107: "CLI: Displaying Current Port Security Settings"
port-security	page 108: "CLI: Configuring Port Security"
<[ethernet] port-list>	page 108: "CLI: Configuring Port Security"
[learn-mode continuous]	page 109: "Adding an Authorized Device to a Port"
[learn-mode static]	page 109: "Adding an Authorized Device to a Port"
[address-limit]	page 109: "Adding an Authorized Device to a Port"
[mac-address]	page 109: "Adding an Authorized Device to a Port"
[action]	page 109: "Adding an Authorized Device to a Port"
no port-security	page 111: "Removing a Device From the "Authorized" List for a Port"
[clear-intrusion-flag]	page 116: "CLI: Checking for Intrusions, Listing Intrusion Alerts, and Resetting Alert Flags"

This section describes the CLI port security command and how the switch acquires and maintains authorized addresses.

Note

Use the global configuration level to execute port-security configuration commands.

Table 4. Port Security Parameters

Parameter	Description
Port List	<code><[ethernet] port-list></code> Identifies the port or ports on which to apply a port security command.
Learn Mode	<p>learn-mode <static continuous> Specifies how the port acquires authorized addresses.</p> <p>Continuous (Default): Appears in the factory-default setting or when you execute no port-security. Allows the port to learn addresses from inbound traffic from any device(s) to which it is connected. In this state, the port accepts traffic from any device(s) to which it is connected. Addresses learned this way appear in the switch and port address tables and age out according to the MAC Age Interval in the System Information configuration screen of the Menu interface or the show system-information listing.</p> <p>Static: Enables you to use the mac-address parameter to specify the MAC addresses of the devices authorized for a port, and the address-limit parameter to specify the number of MAC addresses authorized for the port. You can authorize specific devices for the port, while still allowing the port to accept other, non-specified devices until the device limit has been reached. That is, if you enter fewer MAC addresses than you authorized, the port authorizes the remaining addresses in the order in which it automatically learns them. For example, if you use address-limit to specify three authorized devices, but use mac-address to specify only one authorized MAC address, the port adds the one specifically authorized MAC address to its authorized-devices list and the first two additional MAC addresses it detects. If, for example:</p> <ul style="list-style-type: none"> – You use mac-address to authorize MAC address 0060b0-880a80 for port A4. – You use address-limit to allow three devices on port A4 and the port detects a series of MAC addresses in the following order: <ul style="list-style-type: none"> 1. 080090-1362f2 3. 080071-0c45a1 2. 00f031-423fc1 4. 0060b0-880a80 (the address you authorized with the mac-address parameter) <p>In the above case, port A4 would assume the following list of authorized addresses:</p> <ul style="list-style-type: none"> 080090-1362f2 (the first address the port detected) 00f031-423fc1 (the second address the port detected) 0060b0-880a80 (the address you authorized with the mac-address parameter) <p>The remaining MAC address the port detects, 080071-0c45a1, is not allowed, and is handled as an intruder.</p> <p>Retention of Static Addresses</p> <p>Learned Addresses: In the following two cases, a port in Static learn mode retains a learned MAC address even if you later reboot the switch or disable port security for that port:</p> <ul style="list-style-type: none"> • 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 file (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. <p>To remove an address learned using either of the preceding methods, do one of the following:</p> <ul style="list-style-type: none"> • Delete the address by using no port-security <port-number> mac-address <mac-addr>. • Download a configuration file that does not include the unwanted MAC address assignment. • Reset the switch to its factory-default configuration. <p>Assigned/Authorized Addresses: If you manually assign a MAC address (using port-security <port-number> address-list <mac-addr>) and then execute write memory, the assigned MAC address remains in memory until you do one of the following:</p> <ul style="list-style-type: none"> • Delete it by using no port-security <port-number> mac-address <mac-addr>. • Download a configuration file that does not include the unwanted MAC address assignment. • Reset the switch to its factory-default configuration. <p>Caution: When you use static with a device limit greater than the number of MAC addresses you specify with mac-address, an unwanted device can become “authorized”. This can occur because the port, in order to fulfill the number of devices allowed by the address-limit parameter, automatically adds devices it detects until the specified limit is reached.</p>

Enhancements in Release G.04.05
Configuring and Monitoring Port Security

Parameter	Description
Device Limit	address-limit < <i>integer</i> > When Learn Mode is set to Static , specifies how many authorized devices (MAC addresses) to allow. Range: 1 (the default) to 8.
Action	action <none send-alarm send-disable> Specifies whether an SNMP trap is sent to a network management station when Learn Mode is set to static and the port detects an unauthorized device, or when Learn Mode is set to continuous and there is an address change on a port. None (the default): Prevents an SNMP trap from being sent. Send Alarm: Causes the switch to send an SNMP trap to a network management station. Send Alarm and Disable: Available only in the static learn-mode. Causes the switch to send an SNMP trap to a network management station and disable the port. For information on configuring the switch for SNMP management, see chapter 8.
Address List	mac-address < <i>mac-addr</i> > Available for static learn mode. Allows up to eight authorized devices (MAC addresses) per port, depending on the value specified in the address-limit parameter. If you use mac-address with static , but enter fewer devices than you specified in the address-limit field, the port accepts not only your specified devices, but also as many other devices as it takes to reach the device limit. For example, if you specify four devices, but enter only two MAC addresses, the port will accept the first two non-specified devices it detects, along with the two specifically authorized devices.
Clear Intrusion Flag	clear-intrusion-flag Clears the intrusion flag for a specific port. (See “Reading Intrusion Alerts and Resetting Alert Flags” on page 113.)

CLI: Displaying Current Port Security Settings

The CLI uses the same command to provide two types of port security listings:

- All ports on the switch with their Learn Mode and (alarm) Action
- Only the specified ports with their Learn Mode, Address Limit, (alarm) Action, and Authorized Addresses

Using the CLI To Display Port Security Settings.

Syntax: show port-security

show port-security <port number>

show port-security [<port number>-<port number>]. .[,<port number>]

Without port parameters, show port-security displays Operating Control settings for all ports on a switch. For example:

```
HP4108 (config)# show port-security
Port Security
Port Learn Mode | Action
-----+-----
A1 1 Static | Send Alarm, Disable Port
A2 2 Static | Send Alarm, Disable Port
A3 3 Static | Send Alarm
A4 4 Static | Send Alarm
A5 5 Static | Send Alarm
A6 6 Static | Send Alarm
A7 7 Continuous | None
A8 8 Continuous | None
```

Figure 55. Example Port Security Listing (Ports A7 and A8 Show the Default Setting)

With port numbers included in the command, **show port-security** displays Learn Mode, Address Limit, (alarm) Action, and Authorized Addresses for the specified ports on a switch. The following example lists the full port security configuration for a single port:

```
HP4108 (config)# show port-security 3
Port Security
  Port : A3
  Learn Mode : Static           Address Limit : 1
  Action : Send Alarm
  Authorized Addresses
  -----
  00906d-fdcc00
```

Figure 56. Example of the Port Security Configuration Display for a Single Port

The following command example shows the option for entering a range of ports, including a series of non-contiguous ports. Note that no spaces are allowed in the port number portion of the command string:

```
HP4108 (config)# show port-security A1-A3,A6,A8
```

CLI: Configuring Port Security

Using the CLI, you can:

- Configure port security and edit security settings.
- Add or delete devices from the list of authorized addresses for one or more ports.
- Clear the Intrusion flag on specific ports

Syntax:

```
port-security <port-list>
  [learn-mode continuous]
  [learn-mode static]
  [address-limit <integer>]
  [mac-address <mac-addr>] [<mac-addr> ... <mac-addr>]
  [action <none | send-alarm | send-disable>]
  [clear-intrusion-flag]

no port-security <port-list> mac-address <mac-addr> [<mac-addr> ...
  <mac-addr>]
```

For information on the individual control parameters, see the Port Security Parameter tables on pages 105 and 106.

Specifying Authorized Devices and Intrusion Responses. This example configures port A1 to automatically accept the first device (MAC address) it detects as the only authorized device for that port. (The default device limit is 1.) It also configures the port to send an alarm to a network management station and disable itself if an intruder is detected on the port.

```
HP4108(config)# port-security A1 learn-mode static  
action send-disable
```

The next example does the same as the preceding example, except that it specifies a MAC address of 0c0090-123456 as the authorized device instead of allowing the port to automatically assign the first device it detects as an authorized device.

```
HP4108(config)# port-security A1 learn-mode static  
mac-address 0c0090-123456 action send-disable
```

This example configures port A5 to:

- Allow two MAC addresses, 00c100-7fec00 and 0060b0-889e00, as the authorized devices
- Send an alarm to a management station if an intruder is detected on the port

```
HP4108(config)# port-security 5 learn-mode static  
address-limit 2 mac-address 00c100-7fec00 0060b0-  
889e00 action send-alarm
```

If you manually configure authorized devices (MAC addresses) and/or an alarm action on a port, those settings remain unless you either manually change them or the switch is reset to its factory-default configuration. You can “turn off” authorized devices on a port by configuring the port to continuous Learn Mode, but subsequently reconfiguring the port to static Learn Mode restores those authorized devices.

Adding an Authorized Device to a Port. To simply add a device (MAC address) to a port’s existing Authorized Addresses list, enter the port number with the **mac-address** parameter and the device’s MAC address. *This assumes that Learn Mode is set to static and the Authorized Addresses list is not full* (as determined by the current Address Limit value). For example, suppose port A2 allows two authorized devices, but has only one device in its Authorized Address list:

```
HP4108 (config)# show port-security 1
Port Security
  Port : A1
  Learn Mode : Static
  Action : None
  Authorized Addresses
  -----
  0c0090-123456
  Address Limit : 2
```

Although the Address Limit is set to 2, only one device has been authorized for this port. In this case you can add another without having to also increase the Address Limit.

The Address Limit has not been reached.

Figure 57. Example of Adding an Authorized Device to a Port

With the above configuration for port A1, the following command adds the 0c0090-456456 MAC address as the second authorized address.

```
HP4108 (config)# port-security A1 mac-address 0c0090-456456
```

After executing the above command, the security configuration for port A1 would be:

```
HP4108 (config)# show port-sec A1
Port Security
  Port : A1
  Learn Mode : Static
  Action : None
  Authorized Addresses
  -----
  0c0090-123456
  0c0090-456456
  Address Limit : 2
```

The Address Limit has been reached.

Figure 58. Example of Adding a Second Authorized Device to a Port

(The message `Inconsistent` value appears if the new MAC address exceeds the current Address Limit or specifies a device that is already on the list. Note that if you change a port from static to continuous learn mode, the port retains in memory any authorized addresses it had while in static mode. If you subsequently attempt to convert the port back to static mode with the same authorized address(es), the `Inconsistent` value message appears because the port already has the address(es) in its “Authorized” list.)

If you are adding a device (MAC address) to a port on which the Authorized Addresses list is already full (as controlled by the port’s current Address Limit setting), then you must increase the Address Limit in order to add the device, even if you want to replace one device with another. Using the CLI, you can simultaneously increase the limit and add the MAC address with a single command. For example, suppose port A1 allows one authorized device and already has a device listed:

```
HP4108 (config)# show port-security A1
Port Security
  Port : A1
  Learn Mode : Static           Address Limit : 1
  Action : None
  Authorized Addresses
  -----
  0c0090-123456
```

Figure 59. Example of Port Security on Port A1 with an Address Limit of "1"

To add a second authorized device to port A1, execute a **port-security** command for for port A1 that raises the address limit to 2 and specifies the additional device's MAC address. For example:

```
HP4108 (config)# port-security A1 mac-address 0c0090-456456 address-limit 2
```

Removing a Device From the “Authorized” List for a Port. This command option removes unwanted devices (MAC addresses) from the Authorized Addresses list. (An Authorized Address list is available for each port for which Learn Mode is currently set to “Static”. See the “Address List” entry in the table on page 106.)

Caution

When learn mode is set to static, the Address Limit (address-limit) parameter controls how many devices are allowed in the Authorized Addresses (**mac-address**) for a given port. If you remove a MAC address from the Authorized Addresses list without also reducing the Address Limit by 1, the port may subsequently detect and accept as authorized a MAC address that you do not intend to include in your Authorized Address list. Thus, if you use the CLI to remove a device that is no longer authorized, it is recommended that you first reduce the Address Limit (**address-limit**) integer by 1, as shown below. This prevents the possibility of the same device or another unauthorized device on the network from automatically being accepted as “authorized” for that port.

To remove a device (MAC address) from the “Authorized” list and when the current number of devices equals the Address Limit value, you should first reduce the Address Limit value by 1, then remove the unwanted device.

Note

You can reduce the address limit below the number of currently authorized addresses on a port. This enables you to subsequently remove a device from the “Authorized” list without opening the possibility for an unwanted device to automatically become authorized.

For example, suppose port A1 is configured as shown below and you want to remove 0c0090-123456 from the Authorized Address list:

```
HP4108(config)# show port-sec A1
Port Security
  Port : A1
  Learn Mode : Static           Address Limit : 2
  Action : None
  Authorized Addresses
  -----
  0c0090-123456
  0c0090-456456
```

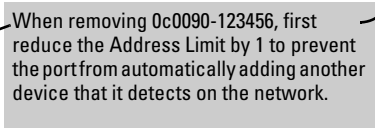


Figure 60. Example of Two Authorized Addresses on Port A1

The following command serves this purpose by removing 0c0090-123456 and reducing the Address Limit to 1:

```
HP4108(config) # port-security A1 address-limit 1
HP4108(config) # no port-security A1 mac-address 0c0090-123456
```

The above command sequence results in the following configuration for port A1:

```
HP2512(config)# show port-sec A1
Port Security
  Port : A1
  Learn Mode : Static           Address Limit : 1
  Action : None
  Authorized Addresses
  -----
  0c0090-456456
```

Figure 61. Example of Port A1 After Removing One MAC Address

Web: Displaying and Configuring Port Security Features

1. Click on the **Security** tab.
2. Click on [Port Security](#).
3. Select the settings you want and, if you are using the Static Learn Mode, add or edit the Authorized Addresses field.
4. Implement your new data by clicking on [Apply Changes](#).

To access the web-based Help provided for the switch, click on [?](#) in the web browser screen.

Reading Intrusion Alerts and Resetting Alert Flags

Notice of Security Violations

When the switch detects an intrusion on a port, it sets an “alert flag” for that port and makes the intrusion information available as described below. *While the switch can detect additional intrusions for the same port, it does not list the next chronological intrusion for that port in the Intrusion Log until the alert flag for that port has been reset.*

When a security violation occurs on a port configured for Port Security, the switch responds in the following ways to notify you:

- The switch sets an alert flag for that port. This flag remains set until:
 - You use either the CLI, menu interface, or web browser interface to reset the flag.
 - The switch is reset to its factory default configuration.
- The switch enables notification of the intrusion through the following means:
 - In the CLI:
 - The `show intrusion-log` command displays the Intrusion Log
 - The `log` command displays the Event Log
 - In the menu interface:
 - The Port Status screen includes a per-port intrusion alert
 - The Event Log includes per-port entries for security violations
 - In the web browser interface:
 - The Alert Log’s Status | Overview window includes entries for per-port security violations
 - The Intrusion Log in the Security | Intrusion Log window lists per-port security violation entries
 - In HP TopTools for Hubs & Switches via an SNMP trap sent to a net management station

How the Intrusion Log Operates

When the switch detects an intrusion attempt on a port, it enters a record of this event in the Intrusion Log. No further intrusion attempts on that port will appear in the Log until you acknowledge the earlier intrusion event by resetting the alert flag.

The Intrusion Log lists the 20 most recently detected security violation attempts, regardless of whether the alert flags for these attempts have been reset. This gives you a history of past intrusion attempts. Thus, for example, if there is an intrusion alert for port A1 and the Intrusion Log shows two or more entries for port 1, only the most recent entry has not been acknowledged (by resetting the alert flag). The other entries give you a history of past intrusions detected on port A1.

Status and Counters - Intrusion Log		
Port	MAC Address	Date / Time
A1	080009-e93d4f	03/07/02 21:09:34
A1	080009-e93d4f	03/07/02 10:18:43

Figure 62. Example of Multiple Intrusion Log Entries for the Same Port

The log shows the most recent intrusion at the top of the listing. You cannot delete Intrusion Log entries (unless you reset the switch to its factory-default configuration). Instead, if the log is filled when the switch detects a new intrusion, the oldest entry is dropped off the listing and the newest entry appears at the top of the listing.

Keeping the Intrusion Log Current by Resetting Alert Flags

When a violation occurs on a port, an alert flag is set for that port and the violation is entered in the Intrusion Log. The switch can detect and handle subsequent intrusions on that port, but will not log another intrusion on the port until you reset the alert flag for either all ports or for the individual port.

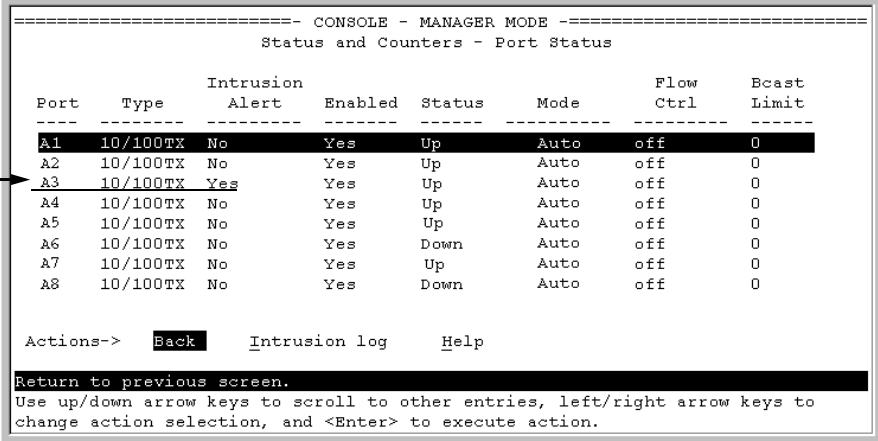
Menu: Checking for Intrusions, Listing Intrusion Alerts, and Resetting Alert Flags

The menu interface indicates per-port intrusions in the Port Status screen, and provides details and the reset function in the Intrusion Log screen.

1. From the Main Menu select:

1. Status and Counters
3. Port Status

The Intrusion Alert column shows "Yes" for any port on which a security violation has been detected.



Port	Type	Intrusion Alert	Enabled	Status	Mode	Flow Ctrl	Bcast Limit
A1	10/100TX	No	Yes	Up	Auto	off	0
A2	10/100TX	No	Yes	Up	Auto	off	0
A3	10/100TX	Yes	Yes	Up	Auto	off	0
A4	10/100TX	No	Yes	Up	Auto	off	0
A5	10/100TX	No	Yes	Up	Auto	off	0
A6	10/100TX	No	Yes	Down	Auto	off	0
A7	10/100TX	No	Yes	Up	Auto	off	0
A8	10/100TX	No	Yes	Down	Auto	off	0

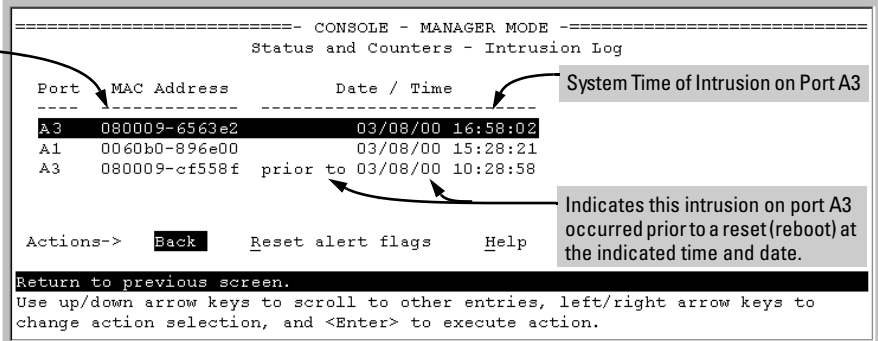
Actions-> **Back** Intrusion log Help

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

Figure 63. Example of Port Status Screen with Intrusion Alert on Port A3

2. Type **I** (Intrusion log) to display the Intrusion Log.

MAC Address of Intruding Device on Port A3



Port	MAC Address	Date / Time
A3	080009-6563e2	03/08/00 16:58:02
A1	0060b0-896e00	03/08/00 15:28:21
A3	080009-cf558f	prior to 03/08/00 10:28:58

Actions-> **Back** Reset alert flags 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.

System Time of Intrusion on Port A3

Indicates this intrusion on port A3 occurred prior to a reset (reboot) at the indicated time and date.

Figure 64. Example of the Intrusion Log Display

The above example shows two intrusions for port A3 and one intrusion for port A1. In this case, only the most recent intrusion at port A3 has not been acknowledged (reset). This is indicated by the following:

- Because the Port Status screen (figure 63 on page 115) does not indicate an intrusion for port A1, the alert flag for the intrusion on port A1 has already been reset.
- Since the switch can show only one uncleared intrusion per port, the older intrusion for port A3 in this example has also been previously reset.

(The intrusion log holds up to 20 intrusion records and deletes an intrusion record only when the log becomes full and a new intrusion is subsequently detected.)

Note also that the “prior to” text in the record for the earliest intrusion means that a switch reset occurred at the indicated time and that the intrusion occurred prior to the reset.

3. To acknowledge the most recent intrusion entry on port A3 and enable the switch to enter a subsequently detected intrusion on this port, type **[R]** (for **Reset alert flags**). (Note that if there are unacknowledged intrusions on two or more ports, this step resets the alert flags for all such ports.)

If you then re-display the port status screen, you will see that the Intrusion Alert entry for port A3 has changed to “**No**”. That is, your evidence that the Intrusion Alert flag has been acknowledged (reset) is that the Intrusion Alert column in the port status display no longer shows “**Yes**” for the port on which the intrusion occurred (port A3 in this example). (Because the Intrusion Log provides a history of the last 20 intrusions detected by the switch, resetting the alert flags does not change its content. Thus, displaying the Intrusion Log again will result in the same display as in figure 64, above.)

CLI: Checking for Intrusions, Listing Intrusion Alerts, and Resetting Alert Flags

The following commands display port status, including whether there are intrusion alerts for any port(s), list the last 20 intrusions, and either reset the alert flag on all ports or for a specific port for which an intrusion was detected. (The record of the intrusion remains in the log. For more information, see “Operating Notes for Port Security” on page 119.)

Syntax:	show interface	List Intrusion Alert status.
	show intrusion-log	List Intrusion Log content.
	clear intrusion-log	Clear Intrusion flags on all ports.
	port-security <port-number>	
	clear-intrusion-flag	Clear Intrusion flag on a specific port.

In the following example, executing **show interface** lists the switch’s port status, which indicates an intrusion alert on port A1.

```

HP4108(config)# show interface
Status and Counters - Port Status

```

Port	Type	Intrusion			Status	Mode	Flow Ctrl	Bcast Limit
		Alert	Enabled					
A1	10/100TX	Yes	Yes	Up	10HDx	off	0	
A2	10/100TX	No	Yes	Up	10HDx	off	0	
A3	10/100TX	No	Yes	Up	10HDx	off	0	

Intrusion Alert on port A1.

Figure 65. Example of an Unacknowledged Intrusion Alert in a Port Status Display

If you wanted to see the details of the intrusion, you would then enter the `show intrusion-log` command. For example:

```

HP4108(config)# show intrusion-log
Status and Counters - Intrusion Log

```

Port	MAC Address	Date / Time
A1	080009-e93d4f	03/07/02 21:09:34
A1	080009-21ae84	03/07/02 17:26:27
A1	080009-e93d4f prior to	03/07/02 17:18:43

MAC Address of latest Intruder on Port A1

Earlier intrusions on port A1 that have already been cleared (that is, the Alert Flag has been reset at least twice before the most recent intrusion occurred).

Dates and Times of Intrusions

Figure 66. Example of the Intrusion Log with Multiple Entries for the Same Port

The above example shows three intrusions for port A1. Since the switch can show only one uncleared intrusion per port, the older two intrusions in this example have already been cleared by earlier use of the `clear intrusion-log` or the `port-security 1 clear-intrusion-flag` command. (The intrusion log holds up to 20 intrusion records, and deletes intrusion records only when the log becomes full and new intrusions are subsequently added.) The “prior to” text in the record for the third intrusion means that a switch reset occurred at the indicated time and that the intrusion occurred prior to the reset.

To clear the intrusion from port A1 and enable the switch to enter any subsequent intrusion for port A1 in the Intrusion Log, execute the `port-security 1 clear-intrusion-flag` command. If you then re-display the port status screen, you will see that the Intrusion Alert entry for port A1 has changed to “No”. That is, your evidence that the Intrusion Alert flag has been reset is the Intrusion Alert column in the port status display no longer shows “Yes” for the port on which the intrusion occurred (port 1 in this example). (Executing `show intrusion-log` again will result in the same display as above.)

Enhancements in Release G.04.05
 Configuring and Monitoring Port Security

```
HP4108(config)# port-security A1 clear-intrusion-flag
HP4108(config)# show interface
```

Status and Counters - Port Status							
Port	Type	Intrusion Alert	Enabled	Status	Mode	Flow Ctrl	Bcast Limit
A1	10/100TX	No	Yes	Up	10HDx	off	0
A2	10/100TX	No	Yes	Up	10HDx	off	0
A3	10/100TX	No	Yes	Up	10HDx	off	0

Intrusion Alert on port A1 is now cleared.

Figure 67. Example of Port Status Screen After Alert Flags Reset

Using the Event Log To Find Intrusion Alerts

The Event Log lists port security intrusions as:

```
W MM/DD/YY HH:MM:SS FFI: port A3 - Security Violation
```

where “W” is the severity level of the log entry and FFI is the system module that generated the entry. For further information, view the Intrusion Log.

From the CLI. Type the log command from the Manager or Configuration level.

Syntax: log <search-text>

For <search-text>, you can use **ffi**, **security**, or **violation**. For example:

Log Listing with Security Violation Detected

```
HP ProCurve Switch 4108GL# log security
Keys:  W=Warning    I=Information
       M=Major      D=Debug
----  Event Log listing: Events Since Boot  ----
W 01/01/02 00:04:30 FFI: port 2 - Security Violation
----  Bottom of Log : Events Listed = 1  ----
```

Log Command with "security" for Search String

Log Listing with No Security Violation Detected

```
HP ProCurve Switch 4108GL# log security
Keys:  W=Warning    I=Information
       M=Major      D=Debug
----  Event Log listing: Events Since Boot  ----
----  Bottom of Log : Events Listed = 0  ----
```

Figure 68. Example of Log Listing With and Without Detected Security Violation

From the Menu Interface: In the Main Menu, click on **4. Event Log** and use **Next page** and **Prev page** to review the Event Log contents.

For More Event Log Information. See “Using the Event Log To Identify Problem Sources” in the "Troubleshooting" chapter of the *Management and Configuration Guide* for your switch.

Web: Checking for Intrusions, Listing Intrusion Alerts, and Resetting Alert Flags

1. Check the Alert Log by clicking on the **Status** tab and the button. If there is a “Security Violation” entry, do the following:
 - a. Click on the **Security** tab.
 - b. Click on . “Ports with Intrusion Flag” indicates any ports for which the alert flag has not been cleared.
 - c. To clear the current alert flags, click on .

To access the web-based Help provided for the switch, click on in the web browser screen.

Operating Notes for Port Security

Identifying the IP Address of an Intruder. The Intrusion Log lists detected intruders by MAC address. If you are using HP TopTools for Hubs & Switches to manage your network, you can use the TopTools inventory reports to link MAC addresses to their corresponding IP addresses. (Inventory reports are organized by device type; hubs, switches, servers, etc.)

Proxy Web Servers. If you are using the switch’s web browser interface through a switch port configured for Static port security, and your browser access is through a proxy web server, then it is necessary to do the following:

- Enter your PC or workstation MAC address in the port’s Authorized Addresses list.
- Enter your PC or workstation’s IP address in the switch’s IP Authorized Managers list. See "Using Authorized IP Managers" in the *Management and Configuration Guide* for your switch.)

Without both of the above configured, the switch detects only the proxy server’s MAC address, and not your PC or workstation MAC address, and interprets your connection as unauthorized.

“Prior To” Entries in the Intrusion Log. If you reset the switch (using the Reset button, Device Reset, or Reboot Switch), the Intrusion Log will list the time of all currently logged intrusions as “prior to” the time of the reset.

Alert Flag Status for Entries Forced Off of the Intrusion Log. If the Intrusion Log is full of entries for which the alert flags have not been reset, a new intrusion will cause the oldest entry to drop off the list, but will not change the alert flag status for the port referenced in the dropped entry. This means that, even if an entry is forced off of the Intrusion Log, no new intrusions can be logged on the port referenced in that entry until you reset the alert flags.

Configuring Rapid Reconfiguration Spanning Tree (RSTP)

This section is related to the information on “Spanning Tree Protocol” in the *Management and Configuration Guide* for your switch, but it primarily describes the new information associated with the new Spanning Tree standard, IEEE 802.1w (RSTP), which is supported by the G.04.05 (or greater) 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 G.04.05 software continues to support.

Overview

RSTP Feature	Default	Menu	CLI	Web
Viewing the RSTP/STP configuration	--	page 130	page 124	n/a
enable/disable RSTP/STP (RSTP is selected as the default protocol)	disabled	page 130	page 125	page 131
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 130	page 126	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 130	page 128	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.

Enhancements in Release G.04.05

Configuring Rapid Reconfiguration Spanning Tree (RSTP)

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 126.

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 129 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 125
protocol-version <rstp stp>	both	page 126
force-version <rstp-operation stp-compatible>	RSTP	page 126
forward-delay <4 - 30>	both	page 126
hello-time <1 - 10>	both	page 126
maximum-age <6 - 40>	both	page 126
priority <0 - 15 0 - 65535>	RSTP STP	page 126
<[ethernet] port-list>	both	page 128
path-cost <1 - 200 000 000>	both	page 128
priority <0 - 15 0 - 65535>	RSTP STP	page 128
edge-port	RSTP	page 128
point-to-point-mac	RSTP	page 128
mcheck	RSTP	page 128
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
-----+-----+-----+-----+-----+-----
A1  10/100TX | 200000 | 8      | Yes  | Force-True     | Yes
A2  10/100TX | 200000 | 8      | Yes  | Force-True     | Yes
A3  10/100TX | 200000 | 8      | Yes  | Force-True     | Yes
A4  10/100TX | 200000 | 8      | Yes  | Force-True     | Yes
A5  10/100TX | 200000 | 8      | Yes  | Force-True     | Yes
A6  10/100TX | 200000 | 8      | Yes  | Force-True     | Yes
A7  10/100TX | 200000 | 8      | Yes  | Force-True     | Yes
A8  10/100TX | 200000 | 8      | Yes  | Force-True     | Yes
A9  10/100TX | 200000 | 8      | Yes  | Force-True     | Yes
A10 10/100TX | 200000 | 8      | Yes  | Force-True     | Yes
A11 10/100TX | 200000 | 8      | Yes  | Force-True     | Yes
A12 10/100TX | 200000 | 8      | Yes  | Force-True     | Yes
- MORE --, next page: Space, next line: Enter, quit: Control-C

```

Figure 69. 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 5. 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 122, 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 128, 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:

```
HP4108 (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 6. 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 129 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
-----+-----+-----+-----+-----+-----+-----
A1    10/100TX | 200000 8       Yes     Force-True     Yes
A2    10/100TX | 200000 8       Yes     Force-True     Yes
A3    10/100TX | 200000 8       Yes     Force-True     Yes
A4    10/100TX | 200000 8       Yes     Force-True     Yes
A5    10/100TX | 200000 8       Yes     Force-True     Yes
A6    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 70. 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 123 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.

Enhancements in Release G.04.05
Configuring Rapid Reconfiguration Spanning Tree (RSTP)

Fast-Uplink Spanning Tree Protocol (STP)

Fast-Uplink STP is an option added to the switch's 802.1d STP to improve the recovery (convergence) time in wiring closet switches with redundant uplinks. Specifically, a Switch 4108GL 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, the 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

Fast-Uplink STP operates only with 802.1d STP and is not available with the Rapid STP (802.1w) feature (page 121).

Caution

In general, fast-uplink spanning tree on the Switch 4108GL 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 Internetworking Protocols*, Addison-Wesley Professional Computing Series, October 1999

Note

When properly implemented, fast-uplink STP offers a method for achieving faster failover times than standard STP, and is intended for this purpose until the true Rapid Convergence STP standard (802.1w) is finalized, approved, and available.

Enhancements in Release G.04.05
Fast-Uplink Spanning Tree Protocol (STP)

To use fast-uplink STP, 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 Switch 4108GL 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:

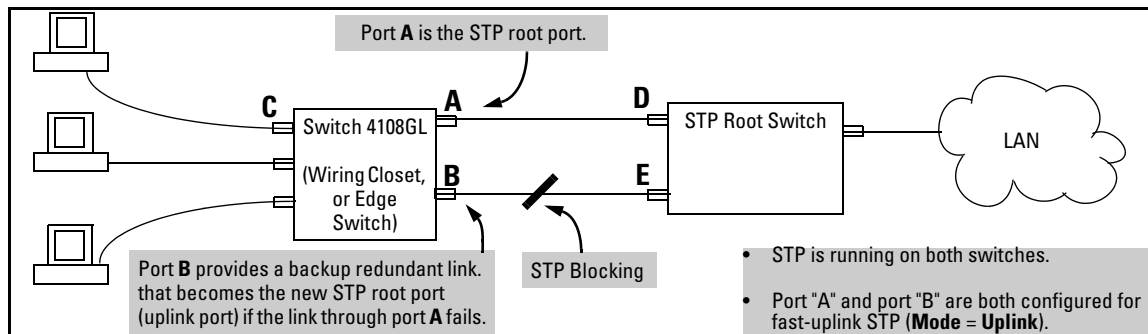


Figure 71. Example of How To Implement Fast-Uplink STP

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 71, 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 72 (page 135) 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 72 (page 135) 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 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 71 on page 134 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 4108GL 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 72, below.)




Figure 72. Example of an Edge Switch in a Topology Configured for STP Fast Uplink

In figure 72, STP is enabled and in its default configuration on all switches, unless otherwise indicated in table 7, below:

Table 7. STP Parameter Settings for Figure 72

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 72.

²This setting ensures that Switch "2" will be the backup root switch for STP in figure 72.

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 72 (page 135) 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 73 is not allowed for fast-uplink operation.




Figure 73. 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 73 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.

1. From the Main Menu select:
 2. **Switch Configuration ...**
 4. **Spanning Tree Operation**
2. 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:

```

===== 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
-----+-----
A3      10/100TX | 200000    8         Yes   Force-True      Yes
A4      10/100TX | 200000    8         Yes   Force-True      Yes
A5      10/100TX | 200000    8         Yes   Force-True      Yes
A6      10/100TX | 200000    8         Yes   Force-True      Yes
A7      10/100TX | 200000    8         Yes   Force-True      Yes
A8      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.

```

- 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 76. In this case, go to step 4 on page 139.

Figure 74. 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 74), do the following:
 - a. Press **[E]** (**E**dit) 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 **S**ave) to save the change and exit from the Spanning Tree Operation screen. you will then see a screen with the following:

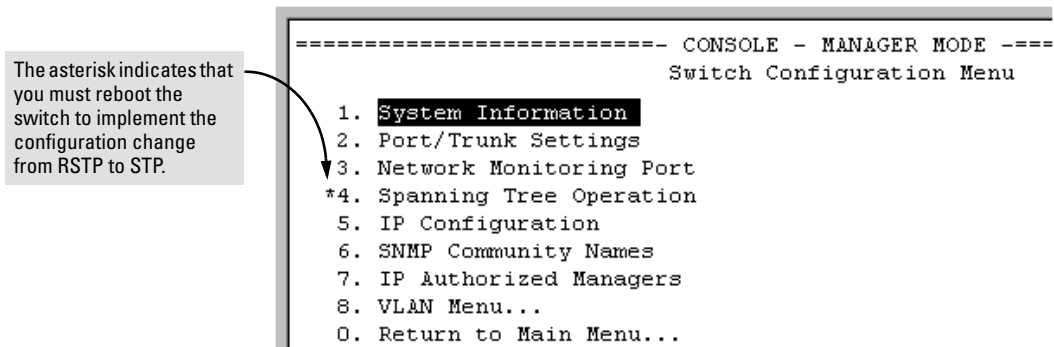


Figure 75. Changing from RSTP to STP Requires a System Reboot

- 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 76).


```

===== 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
A3  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: In the actual menu screen, you must scroll the cursor down the port list to view the trunk configuration (ports A2 and A3).

Figure 76. 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 A1 and Trk1 (using ports A2 and A3) 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 A1.
 - c. Use the Space bar to select **Uplink** as the mode for port A1.
 - 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 is enabled.
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      Uplink
A4     10/100TX | 100    128      Norm
A5     10/100TX | 100    128      Norm
.      .           .      .
.      .           .      .
A24    10/100TX | 100    128      Norm
Trk1   | 100      64      Uplink

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 77. Example of STP Enabled with Two Redundant Links Configured for Fast-Uplink STP

5. Press **[S]** (for **Save**) to save the configuration changes to flash (non-volatile) memory.

To View Fast-Uplink STP Status. Continuing from figures 76 and 77 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 78. Example of STP Status with Trk1 (Trunk 1) as the Path to the STP Root Device

2. Press **[S]** (for **Show 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
-----
A1    10/100TX  10    128    Blocking  0030c1-7fcc40
A4    10/100TX  10    128    Disabled
A5    10/100TX  10    128    Forwarding  0030c1-a914c0
A6    10/100TX  10    128    Forwarding  0030c1-a919c1
.    .
.    .
A24   10/100TX  10    128    Forwarding  0030c1-c884c0
Trk1  Trunk     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 79. Example of STP Port Status with Two Redundant STP Links

In figure 79:

- Port A1 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 A1 blocking (the backup link). (To view the configuration for port A1 and Trk1, see figure 77 on page 140.)

Enhancements in Release G.04.05

Fast-Uplink Spanning Tree Protocol (STP)

- If the link provided by trunk 1 fails (on both ports), then port A1 begins forwarding in fast-uplink STP mode.
- Ports A5, A6, and A24 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 80 and 81 illustrate a possible topology, STP status listing, and STP configuration for a Switch 4108GL with:

- STP enabled and the switch operating as an Edge switch
- Port A1 and trunk 1 (Trk1) configured for fast-uplink STP operation
- Several other ports connected to PC or workstation end nodes

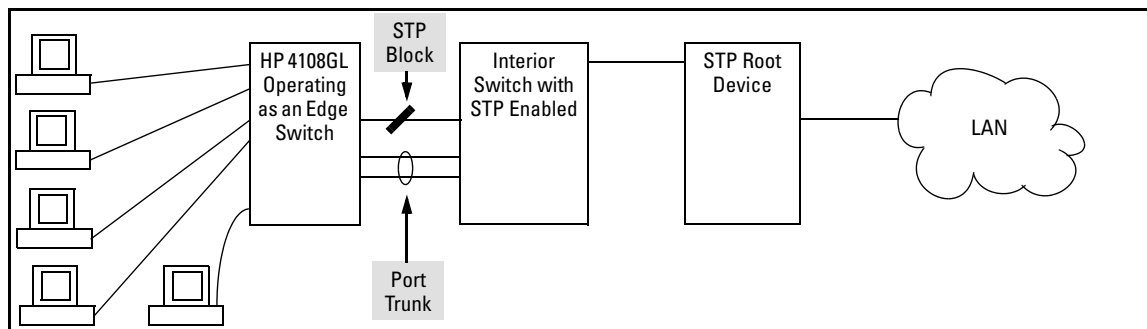


Figure 80. Example Topology for the Listing Shown in Figure 81

```

HP4108 (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
-----+-----
A1    10/100TX   10    128     Blocking   | 0030c1-a9c800
A4    10/100TX   10    128     Disabled   |
A5    10/100TX   10    128     Forwarding  | 0030c1-7fec40
A6    10/100TX   10    128     Forwarding  | 0030c1-a9c800
- MORE --
A7    10/100TX   10    128     Forwarding  | 0030c1-a9c822
A8    10/100TX   10    128     Disabled   |
A9    10/100TX   10    128     Forwarding  | 00a0c9-a234c3
A10   10/100TX   10    128     Forwarding  | 0030c1-449bc0
A11   10/100TX   10    128     Disabled   |
A12   10/100TX   10    128     Disabled   |
Trk1  10/100TX   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 81. Example of a Show Spanning-Tree Listing for the Topology Shown in Figure 80

Enhancements in Release G.04.05
Fast-Uplink Spanning Tree Protocol (STP)

```

HP4108(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
-----+-----
A1 10/100TX | 10 | 128 | Uplink
A4 10/100TX | 10 | 128 | Norm
A5 10/100TX | 10 | 128 | Norm
A6 10/100TX | 10 | 128 | Norm
A7 10/100TX | 10 | 128 | Norm
A8 10/100TX | 10 | 128 | Norm
A9 10/100TX | 10 | 128 | Norm
A10 10/100TX | 10 | 128 | Norm
A11 10/100TX | 10 | 128 | Norm
A12 10/100TX | 10 | 128 | Norm
Trk1 Trunk | 10 | 64 | Uplink
  
```

STP Enabled on the Switch

Fast-Uplink STP Configured on Port 1 and Trunk 1 (Trk1)

Figure 82. Example of a Configuration Supporting the STP Topology Shown in Figure 80

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 80, 81, and 82. (The example assumes that ports A2 and A3 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:

```

HP4108(config)# show spanning-tree
Status and Counters - Spanning Tree Information
Protocol Version : RSTP
STP Enabled : No

Port Type      Cost      Priority State | Designated Bridge
-----+-----
  
```

```

HP4108(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.
HP4108(config)# write mem
HP4108(config)# boot
Device will be rebooted, do you want to continue [y/n]? y
Boot from primary flash
  
```

Lists STP configuration.
Shows the default STP protocol version.

1. Changes the Spanning-Tree protocol to STP (required for Fast-Uplink).
2. Saves the change to the startup-configuration
3. Reboots the switch. (Required for this configuration change.)

Figure 83. 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).

For example:

```
HP4108(config)# spanning-tree e A1,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 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 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.

Listing Switch Configuration and Operation Details for Help in Troubleshooting

Release G.04.05 includes the **show tech** command. 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.

Enhancements in Release G.04.05

Listing Switch Configuration and Operation Details for Help in Troubleshooting

1. In Hyperterminal, click on **Transfer | Capture Text...**




Figure 84. 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 85. 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**:

```
HP4108# 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.

Releases G.03.09, G.03.10, and G.03.13

These three releases did not include enhancements.

Enhancements in Release G.04.05

Listing Switch Configuration and Operation Details for Help in Troubleshooting

Software Fixes

Release G.03.08 was the first software release for the HP Procurve Switch 4108GL.

Release G.03.09	Below
Release G.03.10	Below
Release G.03.13	Page 152
Release G.04.01	Page 152
Release G.04.02	Page 153
Release G.04.03	Page 153
Release G.04.04	Page 153
Release G.04.05	Page 153

Release G.03.09 (Beta Release Only)

Fixed in release G.03.09

- **CDP** — 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.
- **IGMP** — With IGMP enabled, toggling IGMP off and then on again causes all querier intervals to be cut in half.
- **IGMP** — Switch may stop sending IGMP queries on some VLANs.
- **LACP** — Ports are put into Standby mode when they shouldn't be.
- **LACP** — Dynamic LACP creates 2 trunks when it should only create one.
- **LACP** — When an LACP port is put into Standby mode, MAC address learning on the switch may stop.

Release G.03.10

Fixed in release G.03.10

Performance enhancements to the message system, address learning, and SNMP.

Release G.03.13

Fixed in release G.03.13

- **XMODEM** — If the CLI "copy xmodem flash" command is used to download the OS, the switch incorrectly displays one of the two following messages after the validate and write process completes:
 - User timeout, must hit ENTER before starting XMODEM Transfer.or
 - Transfer terminated due to timeout.
- **Inter-module communication problems** — When the switch's MAC address learning function detects the same MAC address on two different modules within a small interval of time (this could happen if two end nodes have the same MAC address or if there is a loop in the network), the address tables on the modules may get out of sync. This can cause module-to-module communication problems for devices connected to the modules whose address tables have become out of synchronization.

Release G.04.01 (Beta Release Only)

Fixed in release G.04.01

- **Crash** — Switch may crash with a message similar to:

```
->Software exception at bcm56xxDmaPoll.c:342--in 'sal_dpc_hi'  
->Msg loss detected
```
- **CLI** — The crash history is lost after the "reload" command is performed from the CLI.
- **CLI** — The response to an incomplete trunk configuration command did not produce the proper message "Incomplete input: Trunk."
- **Flow Control** — Changing Flow Control setting on a port is not reflected in Auto-negotiation's advertised capability.
- **Menu/Web-Browser Interface** — Display of mirror port configuration is inconsistent between menu and WEB interface.
- **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** — Incorrect font size used in VLAN configuration screen.
- **Web-Browser Interface** — User could input an invalid MAC address, i.e. multicast or broadcast address, in the security policy field.

Release G.04.02 (Beta Release Only)

Fixed in release G.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 G.04.03 (Beta Release Only)

Fixed in release G.04.03

Modification of Lab troubleshooting commands.

Release G.04.04 (Beta Release Only)

Fixed in release G.04.04

Modification of Lab troubleshooting commands.

Release G.04.05

Fixed in release G.04.05

Modification of Lab troubleshooting commands.

Index

Numerics

- 3DES ... 12
- 802.1x
 - See *port-based access control*. ... 65

A

- accounting
 - See *RADIUS*.
- address
 - authorized for port security ... 102
- Adobe Acrobat Reader ... 1
- authorized addresses
 - for port security ... 102

B

- broadcast storm ... 121

C

- caution
 - archive config file ... i, 1
- Class of Service
 - priority settings mapped to downstream devices ... 95
- CLI
 - configuring RSTP ... 124
- configuration
 - download ... 2
 - port security ... 103
 - RADIUS
 - See *RADIUS*.
 - RSTP
 - from the CLI ... 124
 - from the menu ... 130
 - per-port parameters ... 128
 - whole switch parameters ... 126
 - running-config file ... 4
 - saving from CLI ... 4

SSH

- See *SSH*.
- startup-config file ... 4
- configuration file
 - software update caution ... i, 1
- configuring RSTP ... 123

D

- DES ... 12
- documentation, download from web ... 1
- download
 - documentation from web ... 1
 - OS to switch ... 2
 - software from web ... 1
 - TFTP ... 2
 - Xmodem ... 3
- downstream device (QoS)
 - effect of priority settings ... 95

E

- enabling RSTP
 - CLI ... 125
 - menu interface ... 130
 - web browser interface ... 131
- Enabling STP
 - CLI ... 125
- enhancements, G.04.05 ... 5
- event log
 - intrusion alerts ... 118

F

- F.01.08 ... 149
- F.01.09 ... 149
- F.01.10 ... 149
- first 4108GL software release ... 151
- friendly port names
 - See *port names, friendly*. ... 6

Index

G

- G.03.08 ... 151
- G.03.09 ... 151
- G.03.10 ... 151
- G.03.13 ... 152
- G.04.01 ... 152
- G.04.02 ... 153
- G.04.03 ... 153
- G.04.04 ... 153
- G.04.05 ... 153
- G.04.05 enhancements ... 5

I

- inconsistent value, message ... 110
- intrusion alarms
 - entries dropped from log ... 119
 - event log ... 118
 - prior to ... 119
- Intrusion Log
 - prior to ... 116–117
- IP
 - reserved port numbers ... 24
- IP preserve
 - DHCP server ... 91
 - overview ... 91
 - rules, operating ... 91
 - summary of effect ... 93

K

- kill command ... 99

M

- MD5
 - See *RADIUS*. ... 38
- menu interface
 - configuring RSTP ... 130
- message
 - inconsistent value ... 110

O

- OpenSSH ... 12
- operating notes
 - port security ... 119
- optimizing RSTP configuration ... 123

P

- part number ... ii
- password security ... 25
- Path Cost
 - comparison of RSTP and STP ... 129
- Perlman, *Interconnections* ... 133
- port
 - security configuration ... 101
- port names, friendly
 - configuring ... 7
 - displaying ... 8
 - summary ... 6
- port security
 - authorized address definition ... 102
 - basic operation ... 101
 - configuring ... 103
 - configuring in browser interface ... 113, 119
 - event log ... 118
 - notice of security violations ... 113
 - operating notes ... 119
 - overview ... 101
 - prior to ... 119
 - proxy web server ... 119
- port trunk
 - with fast-uplink STP ... 145
- port-based access control
 - authenticate switch ... 66
 - authenticate users ... 66
 - authenticator operation ... 66, 69
 - authenticator, show commands ... 81
 - block traffic ... 65
 - blocking non-802.1x device ... 77
 - CHAP ... 65
 - chap-radius ... 75
 - configuration commands ... 72
 - configuration overview ... 71
 - configuration, displaying ... 81
 - configuring method ... 75
 - counters ... 81
 - EAP ... 65
 - EAPOL ... 69
 - eap-radius ... 75
 - enabling on ports ... 72
 - enabling on switch ... 77
 - event log ... 88
 - features ... 65
 - general setup ... 70
 - GVRP ... 87

- local ... 75
- local username and password ... 65
- MD5 ... 69
- messages ... 87
- operation ... 66
- overview ... 65
- port-security, with 802.1x ... 76
- RADIUS ... 65
- RADIUS host IP address ... 76
- rules of operation ... 69
- show commands ... 81
- show commands, supplicant ... 83
- statistics ... 81
- supplicant operation ... 69
- supplicant operation, switch-port ... 68
- supplicant state ... 83
- supplicant statistics, note ... 83
- supplicant, configuring ... 78
- supplicant, configuring switch port ... 80
- supplicant, enabling ... 79
- switch username and password ... 65
- terminology ... 69
- troubleshooting ... 88
- troubleshooting, gvrp ... 84
- used with port-security ... 76
- VLAN operation ... 84
- port-based priority
 - 802.1q VLAN tagging ... 94
 - configuring ... 96
 - messages ... 98
 - overview ... 94
 - priority/queue table ... 95
 - requirement for continuity ... 95
 - rules of operation ... 96
 - troubleshooting ... 98
 - viewing configuration ... 96
- prior to ... 116–117, 119
- Privacy Enhanced Mode (PEM)
 - See *SSH*.
- proxy
 - web server ... 119
- publication data ... ii

R

RADIUS

- accounting ... 37, 50
- accounting, configuration outline ... 52

- accounting, configure server access ... 53
- accounting, configure types on switch ... 54
- accounting, exec ... 51, 54
- accounting, interim updating ... 55
- accounting, network ... 54
- accounting, operating rules ... 51
- accounting, server failure ... 51
- accounting, session-blocking ... 55
- accounting, start-stop method ... 55
- accounting, statistics terms ... 58
- accounting, stop-only method ... 55
- accounting, system ... 51, 54
- authentication options ... 37
- authentication, local ... 48–49
- bypass RADIUS server ... 43
- commands, accounting ... 50
- commands, switch ... 40
- configuration outline ... 41
- configure server access ... 44
- configuring switch global parameters ... 46
- general setup ... 39
- local authentication ... 43
- MD5 ... 38
- messages ... 62
- network accounting ... 50
- operating rules, switch ... 38
- security ... 43
- security note ... 37
- server access order ... 51
- server access order, changing ... 61
- servers, multiple ... 47
- show accounting ... 60
- show authentication ... 59
- SNMP access security not supported ... 37
- statistics, viewing ... 56
- terminology ... 38
- TLS ... 38
- troubleshooting ... 63
- web-browser access controls ... 49
- web-browser security not supported ... 37, 49
- RADIUS accounting
 - See *RADIUS*.
- remote session, terminate ... 99
- reserved port numbers ... 24
- RSTP
 - configuring ... 123
 - configuring per-port parameters ... 128
 - configuring whole switch parameters ... 126

Index

- configuring with the CLI ... 124
- configuring with the menu ... 130
- edge-port parameter ... 128
- enabling from CLI ... 125
- enabling from the menu ... 130
- enabling with the web browser ... 131
- mcheck parameter ... 128
- optimizing the configuration ... 123
- path cost compared to STP ... 129
- path-cost parameter ... 128
- point-to-point-mac parameter ... 128
- priority parameter ... 128
- viewing the configuration ... 124

running-config file ... 4

S

security

- per port ... 101

security violations

- notices of ... 113

security, password

- See *SSH*.

show tech ... 147

software

- download from web ... 1
- download OS to switch ... 2
- fixes ... 151

software update, caution ... i, 1

spanning tree

- caution, fast-uplink ... 133
- configuring per-port parameters ... 128
- configuring RSTP ... 123
- configuring whole-switch parameters ... 126
- configuring with the menu ... 130
- enabling RSTP ... 125
- enabling STP ... 125
- fast-uplink terminology ... 134
- fast-uplink, configuring ... 144
- fast-uplink, menu ... 137
- fast-uplink, operating notes ... 145
- fast-uplink, troubleshooting ... 146
- fast-uplink, viewing status, CLI ... 142
- fast-uplink, viewing status, menu ... 140
- fast-uplink, with port trunks ... 145
- path cost issue ... 129
- RSTP edge port parameter ... 128
- RSTP mcheck parameter ... 128

- RSTP path-cost parameter ... 128
- RSTP point-to-point-mac parameter ... 128
- RSTP priority parameter ... 128
- rules, operating, fast-uplink ... 136
- show tech, copy output ... 147
- viewing the configuration ... 124

SSH

- authenticating switch to client ... 12
- authentication, client public key ... 11
- authentication, user password ... 11
- caution, restricting access ... 26
- caution, security ... 24
- CLI commands ... 17
- client behavior ... 22–23
- client public-key authentication ... 25, 28
- client public-key, clearing ... 32
- client public-key, creating file ... 29
- client public-key, displaying ... 31
- configuring authentication ... 25
- crypto key ... 19
- disabling ... 19
- enable ... 23
- enabling ... 22
- erase host key pair ... 19
- generate host key pair ... 19
- generating key pairs ... 18
- host key pair ... 19
- key, babble ... 19
- key, fingerprint ... 19
- keys, zeroizing ... 18
- key-size ... 24
- known-host file ... 20, 22
- man-in-the-middle spoofing ... 20, 23
- messages, operating ... 33
- OpenSSH ... 12
- operating rules ... 16
- outbound SSH not secure ... 16
- password security ... 25
- password-only authentication ... 25
- passwords, assigning ... 17
- PEM ... 11, 13, 16, 22
- prerequisites ... 13
- public key ... 13, 21
- public key, displaying ... 21
- reserved IP port numbers ... 24
- security ... 20, 24
- SSHv1 ... 11–12
- SSHv1 compatibility ... 20

- SSHv2 ... 11
 - steps for configuring ... 15
 - supported encryption methods ... 12
 - switch key to client ... 20
 - terminology ... 13
 - troubleshooting ... 34
 - unauthorized access ... 26, 32
 - version ... 11
 - zeroize ... 19
 - zeroizing a key ... 18
- startup-config file ... 4
- STP
 - enabling from the CLI ... 125
- summary of enhancements ... 5

T

- Telnet
 - terminate session, kill command ... 99
- terminate remote session ... 99
- TFTP
 - download ... 2
- TLS
 - See *RADIUS*. ... 38
- troubleshooting, SSH. ... 34

V

- value, inconsistent ... 110
- Viewing
 - spanning tree configuration ... 124
- VLAN
 - 802.1x ... 84
 - 802.1x, ID changes ... 87

W

- warranty ... ii
- web browser interface
 - enabling RSTP ... 131
- web browser interface, for configuring
 - port security ... 119
 - port security ... 113
- web server, proxy ... 119
- web site, HP
 - documentation downloads ... 1
 - software downloads ... 1

X

- Xmodem OS download ... 3
- xmodem OS download ... 2



i n v e n t

© 2001-2002 Hewlett-Packard Company. All rights reserved. Reproduction, adaptation, or translation without prior written permission is prohibited except as allowed under the copyright laws.

HP Part Number: 5990-3021
Edition 2, February 2002

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

