HP Switch Software

Access Security Guide

2530-48G-PoE+
2530-24G-PoE+
2530-48G
2530-24G

Software version YA.15.10
November 2012

HP Networking 2530 Switches

Access Security Guide

Applicable Products

HP Switch 2530-48G-PoE+	(J9772A)
HP Switch 2530-24G-PoE+	(J9773A)
HP Switch 2530-48G	(J9775A)
HP Switch 2530-24G	(J9776A)

Trademark Credits

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

Software Credits and Notices

SSH on HP Switches is based on the OpenSSH soft- ware toolkit. This product includes software developed by the OpenSSH Project for use in the OpenSSH Toolkit. For more information on OpenSSH, visit www.openssh.com.

SSL on HP Switches is based on the OpenSSL software toolkit. This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit. For more information on OpenSSL, visit www.openssl.org.

This product includes cryptographic software written by Eric Young (eay@cryptsoft.com). This product includes software written by Tim Hudson (tjh@cryptsoft.com).

Portions of the software on HP switches are based on the lightweight TCP/IP (lwIP) software toolkit by Adam Dunkels, and are covered by the following notices.

Copyright © 2001-2003 Swedish Institute of Computer Science. All rights reserved. Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

3. The name of the author may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE AUTHOR ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This product includes software written by Adam Dunkels (adam@sics.se).

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.

The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

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

Software End User License Agreement and Hardware Limited Warranty

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

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.

Contents

Product Documentation

1 Security Overview

3 Web and MAC Authentication

5 RADIUS Authentication, Authorization and Accounting

6 Configuring Secure Shell (SSH)

7 Configuring Secure Socket Layer (SSL)

8 IPv4 Access Control Lists (ACLs)

9 Traffic/Security Filters and Monitors

10 Configuring Port-Based and Client-Based Access Control (802.1X)

Product Documentation

About Your Switch Manual Set

Electronic Publications

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

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

Security Overview

Introduction

This chapter provides an overview of the security features included on your switch. For detailed information on individual features, see the references provided.

Before you connect your switch to a network, HP strongly recommends that you review the section titled "Getting Started with Access Security" on page 1-6. It outlines potential threats for unauthorized switch and network access, and provides guidelines on how to prepare the switch for secure network operation.

Note   For an introduction to the standard conventions used in this guide, refer to the *Getting Started* chapter in the *Basic Operation Guide* for your switch.

For More Information

For IPv6-specific security settings and features, refer to the *IPV6 Configuration Guide* for your switch.

For the latest version of all HP switch documentation, including Release Notes covering recently added features and other software topics, visit the HP Networking web site at **www.hp.com/Networking/support**.

Access Security Features

This section provides an overview of the switch's access security features, authentication protocols, and methods. Table 9-1 lists these features and provides summary configuration guidelines. For more in-depth information, see the references provided (all chapter and page references are to this *Access Security Guide* unless a different manual name is indicated).

Note   The Management Interface wizard provides a convenient step-by-step method to prepare the switch for secure network operation. See "Quick Start: Using the Management Interface Wizard" on page 1-7 for details.

Table 9-1.   Access Security and Switch Authentication Features

Feature	Default Setting	Security Guidelines	More Information and Configuration Details
Manager password	no password	Configuring a local Manager password is a fundamental step in reducing the possibility of unauthorized access through the switch's Web browser and console (CLI and Menu) interfaces. The Manager password can easily be set by any one of the following methods: • CLI: password manager command, or Management interface wizard • Web browser interface: the password options under the Security tab, or Management interface wizard • Menu interface: Console Passwords option • SNMP	*"Configuring Local Password Security" on page 2-3* *"Quick Start: Using the Management Interface Wizard" on page 1-7* *"Using SNMP To View and Configure Switch Authentication Features" on page 5-17*

Feature	Default Setting	Security Guidelines	More Information and Configuration Details
Telnet and Web-browser access	enabled	The default remote management protocols enabled on the switch are plain text protocols, which transfer passwords in open or plain text that is easily captured. To reduce the chances of unauthorized users capturing your passwords, secure and encrypted protocols such as SSH and SSL (see below for details) should be used for remote access. This enables you to employ increased access security while still retaining remote client access. Also, access security on the switch is incomplete without disabling Telnet and the standard Web browser access. Among the methods for blocking unauthorized access attempts using Telnet or the Web browser are the following two CLI commands: • **no telnet-server**: This command blocks inbound Telnet access. • **no web-management**: This command prevents use of the Web browser interface through http (port 80) server access. If you choose not to disable Telnet and Web browser access, you may want to consider using RADIUS accounting to maintain a record of password-protected access to the switch.	*"Quick Start: Using the Management Interface Wizard" on page 1-7* For more on Telnet and web browser access, refer to the chapter on *"Interface Access and System Information"* in the *Basic Operation Guide.* For RADIUS accounting, refer to *Chapter 5 "RADIUS Authentication, Authorization and Accounting"*
SSH	disabled	SSH provides Telnet-like functions through encrypted, authenticated transactions of the following types: • client public-key authentication: 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. • switch SSH and user password authentication: this option is a subset of the client public-key authentication, and is used if the switch has SSH enabled without a login access configured to authenticate the client's key. In this case, the switch authenticates itself to clients, and users on SSH clients then authenticate themselves to the switch by providing passwords stored on a RADIUS or TACACS+ server, or locally on the switch. • secure copy (SC) and secure FTP (SFTP): By opening a secure, encrypted SSH session, you can take advantage of SC and SFTP to provide a secure alternative to TFTP for transferring sensitive switch information. For more on SC and SFTP, refer to the section titled "Using Secure Copy and SFTP" in the "File Transfers" appendix of the *Management and Configuration Guide* for your switch.	*"Quick Start: Using the Management Interface Wizard" on page 1-7* *Chapter 6" "Configuring Secure Shell (SSH)"*

Feature	Default Setting	Security Guidelines	More Information and Configuration Details
SSL	disabled	Secure Socket Layer (SSL) and Transport Layer Security (TLS) provide remote Web browser access to the switch via authenticated transactions and encrypted paths between the switch and management station clients capable of SSL/TLS operation. The authenticated type includes server certificate authentication with user password authentication.	*"Quick Start: Using the Management Interface Wizard" on page 1-7* *Chapter 7"Configuring Secure Socket Layer (SSL)"*
SNMP	public, unrestricted	In the default configuration, the switch is open to access by management stations running SNMP management applications capable of viewing and changing the settings and status data in the switch's MIB (Management Information Base). Thus, controlling SNMP access to the switch and preventing unauthorized SNMP access should be a key element of your network security strategy.	*"SNMP Security Guidelines" on page 1-10* *"Quick Start: Using the Management Interface Wizard" on page 1-7* *Management and Configuration Guide, Chapter 6,* refer to the section *"Using SNMP Tools To Manage the Switch"*
Authorized IP Managers	none	This feature uses IP addresses and masks to determine whether to allow management access to the switch across the network through the following : • Telnet and other terminal emulation applications • The switch's Web browser interface • SNMP (with a correct community name)	*Chapter 12"Using Authorized IP Managers"*
Secure Management VLAN	disabled	This feature creates an isolated network for managing the HP switches that offer this feature. When a secure management VLAN is enabled, CLI, Menu interface, and Web browser interface access is restricted to ports configured as members of the VLAN.	*Advanced Traffic Management Guide,* refer to the chapter *"Static Virtual LANs (VLANs)"*
ACLs for Management Access Protection	none	ACLs can also be configured to protect management access by blocking inbound IP traffic that has the switch itself as the destination IP address.	*Chapter 8, "IPv4 Access Control Lists (ACLs)"*
TACACS+ Authentication	disabled	This application uses a central server to allow or deny access to TACACS-aware devices in your network. TACACS+ uses username/password sets with associated privilege levels to grant or deny access through either the switch's serial (console) port or remotely, with Telnet. If the switch fails to connect to a TACACS+ server for the necessary authentication service, it defaults to its own locally configured passwords for authentication control. TACACS+ allows both login (read-only) and enable (read/write) privilege level access.	*Chapter 4"TACACS+ Authentication"*

Feature	Default Setting	Security Guidelines	More Information and Configuration Details
RADIUS Authentication	disabled	For each authorized client, RADIUS can be used to authenticate operator or manager access privileges on the switch via the serial port (CLI and Menu interface), Telnet, SSH, and Secure FTP/Secure Copy (SFTP/SCP) access methods.	*Chapter 5"RADIUS Authentication, Authorization and Accounting"*
802.1X Access Control	none	This feature provides port-based or user-based authentication through a RADIUS server to protect the switch from unauthorized access and to enable the use of RADIUS-based user profiles to control client access to network services. Included in the general features are the following: • user-based access control supporting up to 2 authenticated clients per port • port-based access control allowing authentication by a single client to open the port • switch operation as a supplicant for point-to-point connections to other 802.1X-compliant HP switches	*Chapter 10"Configuring Port-Based and Client-Based Access Control (802.1X)"*
Web and MAC Authentication	none	These options are designed for application on the edge of a network to provide port-based security measures for protecting private networks and the switch itself from unauthorized access. Because neither method requires clients to run any special supplicant software, both are suitable for legacy systems and temporary access situations where introducing supplicant software is not an attractive option. Both methods rely on using a RADIUS server for authentication. This simplifies access security management by allowing you to control access from a master database in a single server. It also means the same credentials can be used for authentication, regardless of which switch or switch port is the current access point into the LAN. Web authentication uses a web page login to authenticate users for access to the network. MAC authentication grants access to a secure network by authenticating device MAC addresses for access to the network.	*Chapter 3"Web and MAC Authentication"*

Getting Started with Access Security

HP switches are designed as "plug and play" devices, allowing quick and easy installation in your network. In its default configuration the switch is open to unauthorized access of various types. When preparing the switch for network operation, therefore, HP strongly recommends that you enforce a security policy to help ensure that the ease in getting started is not used by unauthorized persons as an opportunity for access and possible malicious actions.

Since security incidents can originate with sources inside as well as outside of an organization, your access security provisions must protect against internal and external threats while preserving the necessary network access for authorized clients and users. It is important to evaluate the level of management access vulnerability existing in your network and take steps to ensure that all reasonable security precautions are in place. This includes both configurable security options and physical access to the switch.

Switch management access is available through the following methods:

■   Front panel access to the console serial port (see "Physical Security")

■   Inbound Telnet access

■   Web-browser access

■   SNMP access

For guidelines on locking down your switch for remote management access, see "Quick Start: Using the Management Interface Wizard" on page 1-7.

Physical Security

Physical access to the switch allows the following:

■   use of the console serial port (CLI and Menu interface) for viewing and changing the current configuration and for reading status, statistics, and log messages.

■   use of the switch's Clear and Reset buttons for these actions:
   •   clearing (removing) local password protection
   •   rebooting the switch
   •   restoring the switch to the factory default configuration (and erasing any non-default configuration settings)

Keeping the switch in a locked wiring closet or other secure space helps to prevent unauthorized physical access.

As additional precautions, you can do the following:

■   Disable or re-enable the password-clearing function of the Clear button.

■   Configure the Clear button to reboot the switch after clearing any local usernames and passwords.

■   Modify the operation of the Reset+Clear button combination so that the switch reboots, but does not restore the switch's factory default settings.

■   Disable or re-enable password recovery.

For the commands used to configure the Clear and Reset buttons, refer to "Front-Panel Security" on page 2-23.

Quick Start: Using the Management Interface Wizard

The Management Interface wizard provides a convenient step-by-step method to prepare the switch for secure network operation. It guides you through the process of locking down the following switch operations or protocols:

■   setting local passwords

■   restricting SNMP access

■   enabling/disabling Telnet

■   enabling/disabling SSH

■   enabling/disabling remote Web management

■   restricting web access to SSL

■   enabling/disabling USB autorun

■   setting timeouts for SSH/Telnet sessions

The wizard can also be used to view the pre-configured defaults and see the current settings for switch access security. The wizard can be launched either via the CLI (see page 1-8) or the Web browser interface (see page 1-9).

Note   The wizard's security settings can also be configured using standard commands via the CLI, Menu or Web browser interfaces. For full details on preparing and configuring the switch for SSH and SSL operation, refer to chapters 6 and 7 respectively.

CLI: Management Interface Wizard

To configure security settings using the CLI wizard, follow the steps below:

1. At the command prompt, type **setup mgmt-interfaces**.

   The welcome banner appears and the first setup option is displayed
   (**Operator password**). As you advance through the wizard, each setup
   option displays the current value in brackets **[ ]** as shown in Figure 1-1.

```
                    Welcome to the Management Interface Setup Wizard

         This wizard will help you with the initial setup of the various
         management interfaces. The current values are shown in brack-
         ets[]. Type in a new value, or press <Enter> to keep the cur-
         rent value. Press CTRL-C at any time to quit the wizard without
         saving any changes. Press ? for help.
```

Current values are
shown in brackets
(Password entries
must be entered
twice and will
appear as asterisks.)

Type in a new
value to change a
setting, or press
<Enter> to keep
the current value.

```
         Operator password                      [not configured]:
                                          Confirm password:
         Manager password                        [*******]:
                                          Confirm password:
         Restrict SNMP access to SNMPv3 only      [no]:
         SNMPv2 community name              [notpublic]:
         SNMPv2 Community access level  [unrestricted]:
         Telnet enabled                            [yes]:
         SSH enabled                                [no]:
         Web management enabled                    [yes]:
         Restrict Web access to SSL                [no]:
         Timeout for ssh/telnet sessions           [0]:
```

Summary of current
settings (displayed
after last wizard
option has been set)

```
         Operator password                          :
         Manager password                           :*******
         Restrict SNMP access to SNMPv3 only        :no
         SNMPv2 community name                      :notpublic
         SNMPv2 Community access level              :unrestricted
         Telnet enabled                             :yes
         SSH enabled                                :no
         Web management enabled                     :yes
         Restrict Web access to SSL                 :no
         Timeout for ssh/telnet sessions            :0

         Do you want to save these changes?     [yes]:
```

To save these settings, press **[Enter]**.
To cancel any changes, type **[n]** (for no), then press **[Enter].**

Figure 1-1. Example of Management Interface Wizard Configuration

2. When you enter the wizard, you have the following options:

   - To update a setting, type in a new value, or press [**Enter**] to keep the current value.
   - To quit the wizard without saving changes, press [**CTRL-C**] at any time.
   - To access online Help for any option, press [**?**].

   After you have gone through each setup option, the wizard displays the summary configuration together with a prompt to save the changes (see Figure 1-1 on page 1-8 for an example).

3. When the message appears asking if you want to save these changes, you have the following options:

   - To save your changes, press [**Enter**].
   - To cancel any changes without saving, type **[n]** and then press [**Enter**].

   After pressing [**Enter**], the wizard exits to the command line prompt.

CLI Wizard: Operating Notes and Restrictions.

■ Once a password has been configured on the switch, you cannot remove it using the CLI wizard. Passwords can be removed by executing the **no password** command directly from the CLI.

■ When you restrict SNMP access to SNMPv3 only, the options SNMPv2 community name and access level will not appear.

■ The wizard displays the first available SNMPv2 community and allows the user to modify the first community access parameters.

■ The wizard creates a new SNMP community only when no communities have been configured on the switch.

■ The USB Autorun feature is disabled as soon as an operator or manager password is set on the switch. Once a password has been set, the USB autorun option is no longer provided as part of the wizard.

Web: Management Interface Wizard

To use the Management Interface wizard from the Web browser interface, follow the steps below:

1. Click the **Security** tab.

2. Click the **Setup Wizard** button. The Welcome window appears.

   This page allows you to choose between two setup types:

- **Typical**—provides a multiple page, step-by-step method to configure security settings, with on-screen instructions for each option.
- **Advanced**—provides a single summary screen in which to configure all security settings at once.

3. To enter the wizard, choose a setup option and then click **Continue**.

- **Typical** (multi-page setup): when you select this option, you will get an alert indicating that configuration changes will be lost if you click on the Web browser's navigation tabs. Click **OK** to close the alert and then advance through the following setup pages: **Operator Password**, **Manager Password**, **SNMP**, **Telnet**, **SSH**, **Web Management GUI**, **USB Autorun**, **Timeout** (see pages for details and setup options).

  At each page, you have the following options:
  – Enter a new value and/or make a new selection, or click **Continue** to keep the current value and proceed to the next page setup. When you have gone through all configurable options, the summary setup page appears.
  – To quit the Setup and return to the Welcome window without saving any changes, click **Exit** at any time.
  – To return to the previous screen(s), click **Back** at any time.

- **Advanced** (single page setup): when you select this option, the summary setup page appears immediately.

Refer to the WebAgent Online Help for detailed information about using the Management Interface wizard.

SNMP Security Guidelines

In the default configuration, the switch is open to access by management stations running SNMP (Simple Network Management Protocol) management applications capable of viewing and changing the settings and status data in the switch's MIB (Management Information Base). Thus, controlling SNMP access to the switch and preventing unauthorized SNMP access should be a key element of your network security strategy.

General SNMP Access to the Switch. The switch supports SNMP versions 1, 2c, and 3, including SNMP community and trap configuration. The default configuration supports versions 1 and 2c compatibility, which uses plain text and does not provide security options.

HP recommends that you enable SNMP version 3 for improved security. SNMPv3 includes the ability to configure restricted access and to block all non-version 3 messages (which blocks version 1 and 2c unprotected operation).

SNMPv3 security options include:

■ configuring device communities as a means for excluding management access by unauthorized stations

■ configuring for access authentication and privacy

■ reporting events to the switch CLI and to SNMP trap receivers

■ restricting non-SNMPv3 agents to either read-only access or no access

■ co-existing with SNMPv1 and v2c if necessary

SNMP Access to the Authentication Configuration MIB. A management station running an SNMP networked device management application, such as IMC or HP OpenView, can access the switch's management information base (MIB) for read access to the switch's status and read/write access to the switch's authentication configuration (hpSwitchAuth). This means that the switch's default configuration now allows SNMP access to security settings in hpSwitchAuth.

Note on SNMP Access to Authentication MIB

Downloading software and booting for the first time enables SNMP access to the authentication configuration MIB (the default action). If SNMPv3 and other security safeguards are not in place, the switch's authentication configuration MIB is exposed to unprotected SNMP access and you should use the command shown below to disable this access.

If SNMP access to the hpSwitchAuth MIB is considered a security risk in your network, then you should implement the following security precautions when downloading and booting from software release YA.15.XX or greater:

■ If SNMP access to the authentication configuration (hpSwitchAuth) MIB described above is not desirable for your network, then immediately after downloading and booting from the YA.15.XX or greater software for the first time, use the following command to disable this feature:

      snmp-server mib hpswitchauthmib excluded

■ If you choose to leave the authentication configuration MIB accessible, then you should do the following to help ensure that unauthorized workstations cannot use SNMP tools to access the MIB:

a. Configure SNMP version 3 management and access security on the switch.

b. Disable SNMP version 2c on the switch.

For details on this feature, refer to the section titled "Using SNMP To View and Configure Switch Authentication Features" on page 5-17.

For more information on configuring SNMP, refer to the section *"Using SNMP Tools To Manage the Switch"* in the chapter *"Configuring for Network Management Applications"* in the *Management and Configuration Guide* for your switch.

HP PMC Identity-Driven Manager (HP PMC IDM)

HP PMC IDM is a plug-in to IMC and uses RADIUS-based technologies to create a user-centric approach to network access management and network activity tracking and monitoring. HP PMC IDM enables control of access security policy from a central management server, with policy enforcement to the network edge, and protection against both external and internal threats.

Using HP PMC IDM, a system administrator can configure automatic and dynamic security to operate at the network edge when a user connects to the network. This operation enables the network to:

■ approve or deny access at the edge of the network instead of in the core;

■ distinguish among different users and what each is authorized to do;

■ configure guest access without compromising internal security.

Criteria for enforcing RADIUS-based security for HP PMC IDM applications includes classifiers such as:

■ authorized user identity

■ authorized device identity (MAC address)

■ software running on the device

■ physical location in the network

■ time of day

Responses can be configured to support the networking requirements, user (SNMP) community, service needs, and access security level for a given client and device.

For more information on HP PMC IDM, go to the HP Web site at **www.hp.com/ Networking/support**, search on IDM.

Configuring Username and Password Security

Feature	Default	Menu	CLI	Web
Set Usernames	none	—	—	page 2-5
Set a Password	none	page 2-3	page 2-4	page 2-5
Delete Password Protection	n/a	page 2-4	page 2-5	page 2-5
Show front-panel-security	n/a	—	page 1-13	—
Front-panel-security		—	page 1-13	—
password-clear	enabled	—	page 1-13	—
reset-on-clear	disabled	—	page 1-14	—
factory-reset	enabled	—	page 1-15	—
password-recovery	enabled	—	page 1-15	—

Console access includes both the menu interface and the CLI. There are two levels of console access: Manager and Operator. For security, you can set a *password pair* (username and password) on each of these levels.

Note

Usernames are optional. Also, in the menu interface, you can configure passwords, but not usernames. To configure usernames, use the CLI or the Web browser interface.

Usernames and passwords for Manager and Operator access can also be configured using SNMP. For more information, refer to "Using SNMP To View and Configure Switch Authentication Features" on page 5-17.

Usernames and passwords for Manager and Operator access can also be configured using the Management Interface Wizard. For more information, refer to "Quick Start: Using the Management Interface Wizard" on page 1-7.

Level	Actions Permitted
Manager:	Access to all console interface areas. *This is the default level.* That is, if a Manager password has *not* been set prior to starting the current console session, then anyone having access to the console can access any area of the console interface.

Level	Actions Permitted
Operator:	Access to the Status and Counters menu, the Event Log, and the CLI*, but no Configuration capabilities.
	On the Operator level, the configuration menus, Download OS, and Reboot Switch options in the Main Menu are not available.

*Allows use of the ping, link-test, show, menu, exit, and logout commands, plus the enable command if you can provide the Manager password.

To configure password security:

1. Set a Manager password pair (and an Operator password pair, if applicable for your system).

2. Exit from the current console session. A Manager password pair will now be needed for full access to the console.

If you do steps 1 and 2, above, then the next time a console session is started for either the menu interface or the CLI, a prompt appears for a password. Assuming you have protected both the Manager and Operator levels, the level of access to the console interface will be determined by which password is entered in response to the prompt.

If you set a Manager password, you may also want to configure the **Inactivity Time** parameter. (Refer to the *Basic Operation Guide* for your switch.) This causes the console session to end after the specified period of inactivity, thus giving you added security against unauthorized console access.

Note

The manager and operator passwords and (optional) usernames control access to the menu interface, CLI, and Web browser interface.

If you configure only a Manager password (with no Operator password), and in a later session the Manager password is not entered correctly in response to a prompt from the switch, then the switch does not allow management access for that session.
Passwords are case-sensitive.

Caution

If the switch has neither a Manager nor an Operator password, anyone having access to the switch through either Telnet, the serial port, or the Web browser interface can access the switch with full manager privileges. Also, if you configure only an Operator password, entering the Operator password enables full manager privileges.

The rest of this section covers how to:

■ Set passwords

■ Delete passwords

■ Recover from a lost password
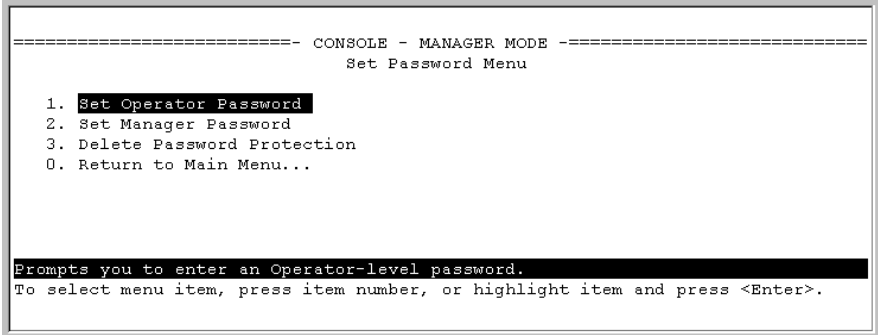
Configuring Local Password Security

Menu: Setting Passwords

As noted earlier in this section, usernames are optional. Configuring a user-name requires either the CLI or the Web browser interface.

1. From the Main Menu select:

   3. Console Passwords

```
==========================- CONSOLE - MANAGER MODE -==============================
                             Set Password Menu

     1. Set Operator Password
     2. Set Manager Password
     3. Delete Password Protection
     0. Return to Main Menu...




Prompts you to enter an Operator-level password.
To select menu item, press item number, or highlight item and press <Enter>.
```

Figure 2-1. The Set Password Screen

2. To set a new password:

   a. Select **Set Manager Password** or **Set Operator Password**. You will then be prompted with **Enter new password**.

   b. Type a password of up to 16 ASCII characters with no spaces and press **[Enter]**. (Remember that passwords are case-sensitive.)

   c. When prompted with **Enter new password again**, retype the new pass-word and press **[Enter]**.

After you configure a password, if you subsequently start a new console session, you will be prompted to enter the password. (If you use the CLI or Web browser interface to configure an optional username, the switch will prompt you for the username, and then the password.)

To Delete Password Protection (Including Recovery from a Lost Password): This procedure deletes *all* usernames (if configured) and passwords (Manager and Operator).

If you have physical access to the switch, press and hold the Clear button (on the front of the switch) for a minimum of one second to clear all password protection, then enter new passwords as described earlier in this chapter.

If you do not have physical access to the switch, you will need Manager-Level access:

1.   Enter the console at the Manager level.

2.   Go to the **Set Passwords** screen as described above.

3.   Select **Delete Password Protection**. You will then see the following prompt:

     Continue Deletion of password protection? No

4.   Press the Space bar to select **Yes**, then press **[Enter]**.

5.   Press **[Enter]** to clear the Password Protection message.

To Recover from a Lost Manager Password: If you cannot start a console session at the Manager level because of a lost Manager password, you can clear the password by getting physical access to the switch and pressing and holding the Clear button for a minimum of one second. This action deletes all passwords and usernames (Manager and Operator) used by both the console and the Web browser interface.

CLI: Setting Passwords and Usernames
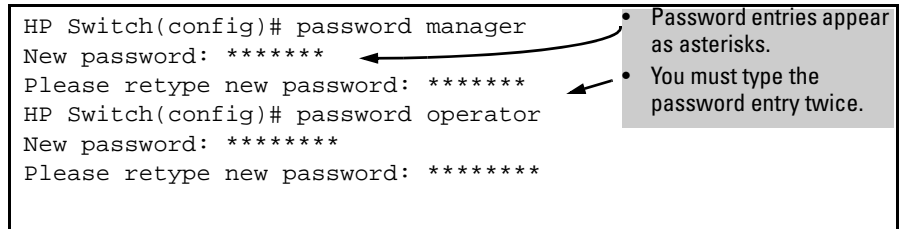
Commands Used in This Section

password	See below.

Configuring Manager and Operator Passwords.

Syntax: [no] password <manager | operator | all > [user-name *ASCII-STR*]
               [no] password

```
HP Switch(config)# password manager
New password: *******
Please retype new password: *******
HP Switch(config)# password operator
New password: ********
Please retype new password: ********
```

- Password entries appear as asterisks.
- You must type the password entry twice.

Figure 2-2. Example of Configuring Manager and Operator Passwords

To Remove Password Protection.   Removing password protection means to eliminate password security. This command prompts you to verify that you want to remove one or both passwords, then clears the indicated password(s). (This command also clears the username associated with a password you are removing.) For example, to remove the Operator password (and username, if assigned) from the switch, you would do the following:

```
HP Switch(config)# no password
Password protection will be deleted, do you want to continue [y/n]? y
HP Switch(config)#
```

Press [Y] (for yes) and press [Enter].
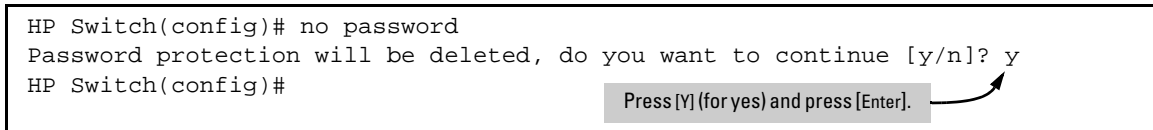
Figure 2-3. Removing a Password and Associated Username from the Switch

The effect of executing the command in figure 2-3 is to remove password protection from the Operator level. (This means that anyone who can access the switch console can gain Operator access without having to enter a user-name or password.)

Web: Setting Passwords and Usernames

In the Web browser interface you can enter passwords and (optional) user-names.

To Configure (or Remove) Usernames and Passwords in the Web Browser Interface.

1.   Click on the **Security** tab.

Click on **[Device Passwords]**.

2.   Do one of the following:

- To set username and password protection, enter the usernames and passwords you want in the appropriate fields.

- • To remove username and password protection, leave the fields blank.

3. Implement the usernames and passwords by clicking on **[Apply Changes]**.

To access the Web-based help provided for the switch, click on **[?]** in the Web browser screen.

Saving Security Credentials in a Config File

You can store and view the following security settings in the running-config file associated with the current software image by entering the **include-credentials** command (formerly this information was stored only in internal flash memory):

■ Local manager and operator passwords and (optional) user names that control access to a management session on the switch through the CLI, menu interface, or web browser interface

■ SNMP security credentials used by network management stations to access a switch, including authentication and privacy passwords

■ Port-access passwords and usernames used as 802.1X authentication credentials for access to the switch

■ TACACS+ encryption keys used to encrypt packets and secure authentication sessions with TACACS+ servers

■ RADIUS shared secret (encryption) keys used to encrypt packets and secure authentication sessions with RADIUS servers

■ Secure Shell (SSH) public keys used to authenticate SSH clients that try to connect to the switch.

Benefits of Saving Security Credentials

The benefits of including and saving security credentials in a configuration file are as follows:

■ After making changes to security parameters in the running configuration, you can experiment with the new configuration and, if necessary, view the new security settings during the session. After verifying the configuration, you can then save it permanently by writing the settings to the startup-config file.

■ By permanently saving a switch's security credentials in a configuration file, you can upload the file to a TFTP server or Xmodem host, and later download the file to the HP switches on which you want to use the same security settings without having to manually configure the settings (except for SNMPv3 user parameters) on each switch.

■ By storing different security settings in different files, you can test different security configurations when you first download a new software version that supports multiple configuration files, by changing the configuration file used when you reboot the switch.

For more information about how to experiment with, upload, download, and use configuration files with different software versions, refer to the following:

■ The chapter on "Switch Memory and Configuration" in the *Basic Operation Guide*.

■ "Configuring Local Password Security" on page 2-3 in this guide.

Enabling the Storage and Display of Security Credentials

To enable the security settings, enter the **include-credentials** command.

Syntax: [no] include-credentials

> *Enables the inclusion and display of the currently configured manager and operator usernames and passwords, RADIUS shared secret keys, SNMP and 802.1X authenticator (port-access) security credentials, and SSH client public-keys in the running configuration. (Earlier software releases store these security configuration settings only in internal flash memory and do not allow you to include and view them in the running-config file.)*
>
> *To view the currently configured security settings in the running configuration, enter one of the following commands:*
>
> ■ **show running-config**: *Displays the configuration settings in the current running-config file.*
>
> ■ **write terminal**: *Displays the configuration settings in the current running-config file.*
>
> *For more information, refer to "Switch Memory and Configuration" in the* Basic Operation Guide.
>
> *The "***no***" form of the command disables only the display and copying of these security parameters from the running configuration, while the security settings remain active in the running configuration.*
>
> *Default: The security credentials described in "Security Settings that Can Be Saved" on page 2-8 are not stored in the running configuration.*

Security Settings that Can Be Saved

The security settings that can be saved to a configuration file are:

■ Local manager and operator passwords and user names

■ SNMP security credentials, including SNMPv1 community names and SNMPv3 usernames, authentication, and privacy settings

■ 802.1X port-access passwords and usernames

■ TACACS+ encryption keys

■ RADIUS shared secret (encryption) keys

■ Public keys of SSH-enabled management stations that are used by the switch to authenticate SSH clients that try to connect to the switch

Local Manager and Operator Passwords

The information saved to the running-config file when the **include-credentials** command is entered includes:

```
password manager [user-name <name>] <hash-type> <pass-hash>
password operator [user-name <name>] <hash-type> <pass-hash>
```

where
<name> is an alphanumeric string for the user name assigned to the manager or operator.
<hash-type> indicates the type of hash algorithm used: SHA-1 or plain text.
<pass-hash> is the SHA-1 authentication protocol's hash of the password or clear ASCII text.

For example, a manager username and password may be stored in a running-config file as follows:

```
  password manager user-name George SHA1
2fd4e1c67a2d28fced849ee1bb76e7391b93eb12
```

Use the **write memory** command to save the password configurations in the startup-config file. The passwords take effect when the switch boots with the software version associated with that configuration file.

C a u t i o n   If a startup configuration file includes other security credentials, but does not contain a manager or operator password, the switch will not have password protection and can be accessed through Telnet, the serial port, or web interface with full manager privileges.

Front-Panel Security

The front-panel security features provide the ability to independently enable or disable some of the functions of the two buttons located on the front of the switch for clearing the password (Clear button) or restoring the switch to its factory default configuration (Reset+Clear buttons together). The ability to disable Password Recovery is also provided for situations which require a higher level of switch security.

The front-panel Security features are designed to prevent malicious users from:

- Resetting the password(s) by pressing the Clear button
- Restoring the factory default configuration by using the Reset+Clear button combination.
- Gaining management access to the switch by having physical access to the switch itself

When Security Is Important

Some customers require a high level of security for information. Also, the Health Insurance Portability and Accountability Act (HIPAA) of 1996 requires that systems handling and transmitting confidential medical records must be secure.

It used to be assumed that only system and network administrators would be able to get access to a network switch because switches were typically placed in secure locations under lock and key. For some customers this is no longer true. Others simply want the added assurance that even if someone did manage to get to the switch that data would still remain secure.

If you do not invoke front-panel security on the switch, user-defined passwords can be deleted by pushing the Clear button on the front panel. This function exists so that if customers forget the defined passwords they can still get back into the switch and reset the passwords. This does, however, leave the switch vulnerable when it is located in an area where non-authorized people have access to it. Passwords could easily be cleared by pressing the Clear button. Someone who has physical access to the switch may be able to erase the passwords (and possibly configure new passwords) and take control of the switch.

As a result of increased security concerns, customers now have the ability to stop someone from removing passwords by disabling the Clear and/or Reset buttons on the front of the switch.

Front-Panel Button Functions

The front panel of the switch includes the Reset button and the Clear button.
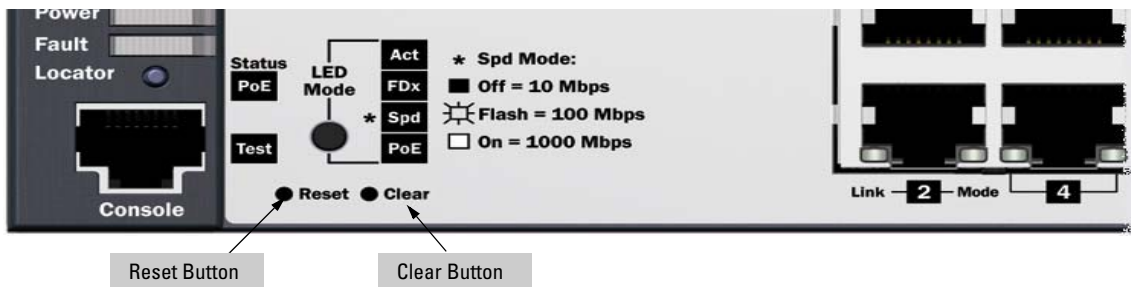
Figure 2-4. Example Front-Panel Button Locations

Clear Button

Pressing the Clear button alone for one second resets the password(s) configured on the switch.
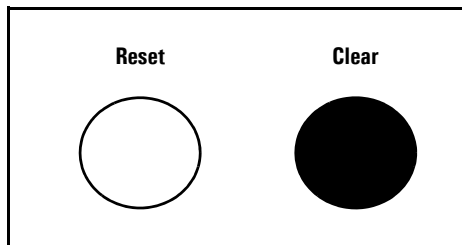
Figure 2-5. Press the Clear Button for One Second To Reset the Password(s)

Reset Button

Pressing the Reset button alone for one second causes the switch to reboot.
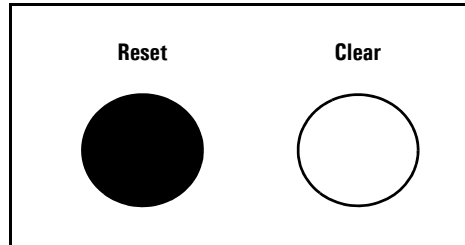
Figure 2-6. Press and hold the Reset Button for One Second To Reboot the Switch

Restoring the Factory Default Configuration

You can also use the Reset button *together* with the Clear button (Reset+Clear) to **restore the factory default configuration** for the switch. To do this:

1. Press and hold the Reset button.
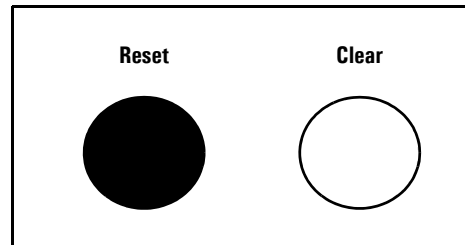
2. While holding the Reset button, press and hold the Clear button.
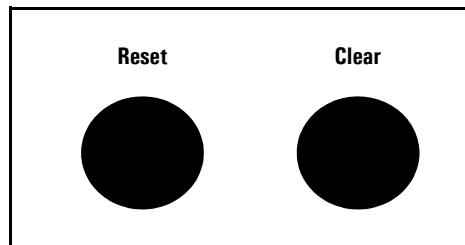
3. Release the Reset button and wait for about one second for the Self-Test LED to start flashing.
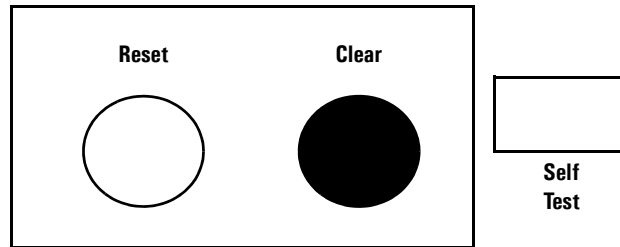
4. When the Self-Test LED begins flashing, release the Clear button

.
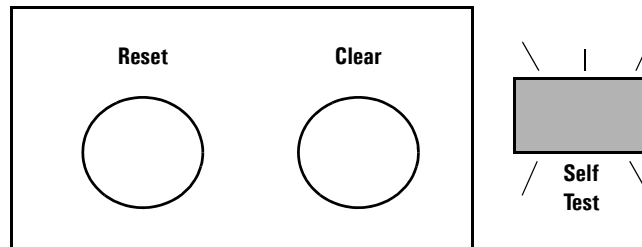
This process restores the switch configuration to the factory default settings.

Configuring Front-Panel Security

Using the **front-panel-security** command from the global configuration context in the CLI you can:

- Disable or re-enable the password-clearing function of the Clear button. Disabling the Clear button means that pressing it does not remove local password protection from the switch. (This action affects the Clear button when used alone, but does not affect the operation of the Reset+Clear combination described under "Restoring the Factory Default Configuration" on page 2-11.)

- Configure the Clear button to reboot the switch after clearing any local usernames and passwords. This provides an immediate, visual means (plus an Event Log message) for verifying that any usernames and passwords in the switch have been cleared.

- Modify the operation of the Reset+Clear combination (page 2-11) so that the switch still reboots, but does *not* restore the switch's factory default configuration settings. (Use of the Reset button alone, to simply reboot the switch, is not affected.)

- Disable or re-enable Password Recovery.

Syntax: show front-panel-security

> *Displays the current front-panel-security settings:*

> **Clear Password:** *Shows the status of the Clear button on the front panel of the switch.* **Enabled** *means that pressing the Clear button erases the local usernames and passwords configured on the switch (and thus removes local password protection from the switch).* **Disabled** *means that pressing the Clear button does not remove the local usernames and passwords configured on the switch. (Default:* **Enabled***.)*

> **Reset-on-clear:** *Shows the status of the reset-on-clear option (***Enabled** *or* **Disabled***). When reset-on-clear is disabled and Clear Password is enabled, then pressing the Clear button erases the local usernames and passwords from the switch. When reset-on-clear is enabled, pressing the Clear button erases the local usernames and passwords from the switch and reboots the switch. (Enabling* **reset-on-clear** *automatically enables* **clear-password***.) (Default:* **Disabled***.)*

> **Factory Reset:** *Shows the status of the Reset button on the front panel of the switch. Enabled means that pressing the Reset button reboots the switch and also enables the Reset button to be used with the Clear button (page 2-11) to reset the switch to its factory-default configuration. (Default:* **Enabled***.)*

> **Password Recovery:** *Shows whether the switch is configured with the ability to recover a lost password. (Refer to "Password Recovery Process" on page 2-19.) (Default:* **Enabled***.)*

> > *CAUTION: Disabling this option removes the ability to recover a password on the switch. Disabling this option is an extreme measure and is not recommended unless you have the most urgent need for high security. If you disable password-recovery and then lose the password, you will have to use the Reset and Clear buttons (page 2-11) to reset the switch to its factory-default configuration and create a new password.*

For example, **show front-panel-security** produces the following output when the switch is configured with the default front-panel security settings.

```
HP Switch(config)# show front-panel-security
Clear Password        - Enabled
  Reset-on-clear      - Disabled
Factory Reset         - Enabled
Password Recovery     - Enabled
```

Figure 2-7. The Default Front-Panel Security Settings

Disabling the Clear Password Function of the Clear Button on the Switch's Front Panel

Syntax: no front-panel-security password-clear

> *In the factory-default configuration, pressing the Clear button on the switch's front panel erases any local usernames and passwords configured on the switch. This command disables the password clear function of the Clear button, so that pressing it has no effect on any local usernames and passwords. (Default: **Enabled**.)*
>
> > ***Note:*** *Although the Clear button does not erase passwords when disabled, you can still use it with the Reset button (*Reset+Clear*) to restore the switch to its factory default configuration, as described under "Restoring the Factory Default Configuration" on page 2-11.*

This command displays a Caution message in the CLI. If you want to proceed with disabling the Clear button, type **[Y]**; otherwise type **[N]**. For example:

```
HP Switch(config)# no front-panel-security password-clear
                       **** CAUTION ****
Disabling the clear button prevents switch passwords from being easily reset or
recovered. Ensure that you are familiar with the front panel security options
before proceeding.

Continue with disabling the clear button [y/n]? y

HP Switch(config)# show front-panel-security
Clear Password        - Disabled
Factory Reset         - Enabled
Password Recovery     - Enabled
```

Indicates the command has disabled the Clear button on the switch's front panel. In this case the Show command does not include the **reset-on-clear** status because it is inoperable while the Clear Password functionality is disabled, and must be reconfigured whenever Clear Password is re-enabled .

Figure 2-8. Example of Disabling the Clear Button and Displaying the New Configuration

Re-Enabling the Clear Button on the Switch's Front Panel and Setting or Changing the "Reset-On-Clear" Operation

Syntax: [no] front-panel-security password-clear reset-on-clear

> *This command does both of the following:*
> - *Re-enables the password-clearing function of the Clear button on the switch's front panel.*
> - *Specifies whether the switch reboots if the Clear button is pressed.*
>
> *To re-enable password-clear, you must also specify whether to enable or disable the* **reset-on-clear** *option.*
> *Defaults:*
> - password-clear: ***Enabled***.
> - reset-on-clear: ***Disabled***.
>
> *Thus:*
> - *To enable password-clear with reset-on-clear disabled, use this syntax:*
>
>   no front-panel-security password-clear reset-on-clear
> - *To enable password-clear with reset-on-clear also enabled, use this syntax:*
>
>   front-panel-security password-clear reset-on-clear
>
> *(Either form of the command <u>enables</u> password-clear.)*
>
> ***Note:*** *If you disable* **password-clear** *and also disable the* **password-recovery** *option, you can still recover from a lost password by using the Reset+Clear button combination at reboot as described on page 2-11. Although the Clear button does not erase passwords when disabled, you can still use it with the Reset button (*Reset+Clear*) to restore the switch to its factory default configuration. You can then get access to the switch to set a new password.*

For example, suppose that **password-clear** is disabled and you want to restore it to its default configuration (enabled, with **reset-on-clear** disabled).

```
HP Switch(config)# show front-panel-security          Shows password-clear disabled.
Clear Password      - Disabled
Factory Reset       - Enabled                          Enables password-clear, with reset-on-
Password Recovery   - Enabled                          clear disabled by the "no" statement at
                                                       the beginning of the command.

HP Switch(config)# no front-panel-security password-clear reset-on-clear

HP Switch(config)# show front-panel-security
Clear Password      - Enabled                          Shows password-clear enabled, with
  Reset-on-clear    - Disabled                         reset-on-clear disabled.
Factory Reset       - Enabled
Password Recovery   - Enabled
```

Figure 2-9.   Example of Re-Enabling the Clear Button's Default Operation

Changing the Operation of the Reset+Clear Combination

In their default configuration, using the Reset+Clear buttons in the combination described under "Restoring the Factory Default Configuration" on page 2-11 replaces the switch's current startup-config file with the factory-default startup-config file, then reboots the switch, and removes local password protection. *This means that anyone who has physical access to the switch could use this button combination to replace the switch's current configuration with the factory-default configuration, and render the switch accessible without the need to input a username or password.* You can use the **factory-reset** command to prevent the Reset+Clear combination from being used for this purpose.

Syntax: [no] front-panel-security factory-reset

> *Disables or re-enables the following functions associated with using the Reset+Clear buttons in the combination described* under "Restoring the Factory Default Configuration" on page 2-11:
>
> • *Replacing the current startup-config file with the factory-default startup-config file*
> • *Clearing any local usernames and passwords configured on the switch*
>
> (***Default:** Both functions enabled.*)
>
> ***Notes:** The Reset+Clear button combination always reboots the switch, regardless of whether the "**no**" form of the command has been used to disable the above two functions. Also, if you disable* **factory-reset**, *you cannot disable the* **password-recovery** *option, and the reverse.*

```
HP Switch(config)# no front-panel-security factory-reset
                          **** CAUTION ****
Disabling the factory reset option prevents switch configuation and passwords
from being easily reset or recovered. Ensure that you are familiar with the
front panel security options before proceeding.

Continue with disabling the factory reset option[y/n]? y

HP Switch(config)# show front-panel-security
Clear Password       - Enabled
  Reset-on-clear     - Disabled
Factory Reset        - Disabled
Password Recovery    - Enabled
```

The command to disable the factory-reset operation produces this caution. To complete the command, press **[Y]**. To abort the command, press **[N]**.

Completes the command to disable the factory reset option.

Displays the current front-panel-security configuration, with Factory Reset disabled.

Figure 2-10. Example of Disabling the Factory Reset Option

Password Recovery

The password recovery feature is enabled by default and provides a method for regaining management access to the switch (without resetting the switch to its factory default configuration) in the event that the system administrator loses the local manager username (if configured) or password. Using Password Recovery requires:

■ **password-recovery** enabled (the default) on the switch prior to an attempt to recover from a lost username/password situation

■ Contacting HP Customer Care Center to acquire a one-time-use password

Disabling or Re-Enabling the Password Recovery Process

Disabling the password recovery process means that the only method for recovering from a lost manager username (if configured) and password is to reset the switch to its factory-default configuration, which removes any non default configuration settings.

C a u t i o n   Disabling **password-recovery** requires that **factory-reset** be enabled, and locks out the ability to recover a lost manager username (if configured) and password on the switch. In this event, there is no way to recover from a lost manager username/password situation without resetting the switch to its factory-default configuration. This can disrupt network operation and make it necessary to temporarily disconnect the switch from the network to prevent unauthorized access and other problems while it is being reconfigured. Also, with **factory-reset** enabled, unauthorized users can use the Reset+Clear button combination to reset the switch to factory-default configuration and gain management access to the switch.

Syntax: [no] front-panel-security password-recovery

> *Enables or (using the "**no**" form of the command) disables the ability to recover a lost password.*
>
> *When this feature is enabled, the switch allows management access through the password recovery process described below. This provides a method for recovering from a lost manager username (if configured) and password. When this feature is disabled, the password recovery process is disabled and the only way to regain management access to the switch is to use the Reset+Clear button combination (page 2-11) to restore the switch to its factory default configuration.*
>
> > *Note: To disable* **password-recovery:**
> > > – *You must have physical access to the front panel of the switch.*
> > > – *The* **factory-reset** *parameter must be enabled (the default).*
>
> *(Default: Enabled.)*

Steps for Disabling Password-Recovery.

1. Set the CLI to the global interface context.

2. Use **show front-panel-security** to determine whether the factory-reset parameter is enabled. If it is disabled, use the **front-panel-security factory-reset** command to enable it.

3. Press and release the Clear button on the front panel of the switch.

4. Within 60-seconds of pressing Clear button, enter the following command:

   no front-panel-security password-recovery

5. Do one of the following after the "**CAUTION**" message appears:
   • If you want to complete the command, press **[Y]** (for "Yes").
   • If you want to abort the command, press **[N]** (for "No")

Figure 2-11 shows an example of disabling the **password-recovery** parameter.

```
HP Switch(config)# no front-panel-security password-recovery
                         **** CAUTION ****
Disabling the clear button without password recovery prevents switch passwords
from being reset. If the switch password is lost, restoring the default factory
configuration will be required to regain access!

Continue with disabling password recovery [y/n]? y
```

Figure 2-11. Example of the Steps for Disabling Password-Recovery

Password Recovery Process

If you have lost the switch's manager username/password, but **password-recovery** is enabled, then you can use the Password Recovery Process to gain management access to the switch with an alternate password supplied by HP.

N o t e   If you have disabled **password-recovery**, which locks out the ability to recover a manager username/password pair on the switch, then the only way to recover from a lost manager username/password pair is to use the Reset+Clear button combination described under "Restoring the Factory Default Configuration" on page 2-11. This can disrupt network operation and make it necessary to temporarily disconnect the switch from the network to prevent unauthorized access and other problems while it is being reconfigured.

To use the **password-recovery** option to recover a lost password:

1. Note the switch's base MAC address. It is shown on the label located on the upper right front corner of the switch.

2. Contact your HP Customer Care Center for further assistance. Using the switch's MAC address, the HP Customer Care Center will generate and provide a "one-time use" alternate password you can use with the to gain management access to the switch. Once you gain access, you can configure a new, known password.

N o t e   The alternate password provided by the HP Customer Care Center is valid only for a single login attempt.

You cannot use the *same "one-time-use" password* if you lose the password a second time. Because the password algorithm is randomized based upon your switch's MAC address, the password will change as soon as you use the *"one-time-use" password* provided to you by the HP Customer Care Center.

Web and MAC Authentication

Overview

Feature	Default	Menu	CLI	Web
Configure Web Authentication	n/a	—	3-16	—
Configure MAC Authentication	n/a	—	3-23	—
Display Web Authentication Status and Configuration	n/a	—	3-27	—
Display MAC Authentication Status and Configuration	n/a	—	3-28	—

Web and MAC Authentication are designed for employment on the "edge" of a network to provide port-based security measures for protecting private networks and the switch itself from unauthorized access. Because neither method requires clients to run any special supplicant software, both are suitable for legacy systems and temporary access situations where introducing supplicant software is not an attractive option. Both methods rely on using a RADIUS server for authentication. This simplifies access security management by allowing you to control access from a master database in a single server. (You can use up to three RADIUS servers to provide backups in case access to the primary server fails.) It also means the same credentials can be used for authentication, regardless of which switch or switch port is the current access point into the LAN.

Web Authentication (Web-Auth)

This method uses a Web page login to authenticate users for access to the network. When a user connects to the switch and opens a Web browser the switch automatically presents a login page. The user then enters a username and password, which the switch forwards to a RADIUS server for authentication. After authentication, the switch grants access to the secured network. Other than a Web browser, the client needs no special supplicant software.

Note   Client Web browsers may not use a proxy server to access the network.

MAC Authentication (MAC-Auth)

This method grants access to a secure network by authenticating devices for access to the network. When a device connects to the switch, either by direct link or through the network, the switch forwards the device's MAC address to the RADIUS server for authentication. The RADIUS server uses the device MAC address as the username and password, and grants or denies network access in the same way that it does for clients capable of interactive logons. (The process does not use either a client device configuration or a logon session.) MAC authentication is well-suited for clients that are not capable of providing interactive logons, such as telephones, printers, and wireless access points. Also, because most RADIUS servers allow for authentication to depend on the source switch and port through which the client connects to the network, you can use MAC-Auth to "lock" a particular device to a specific switch and port.

Concurrent Web and MAC Authentication

Web authentication and MAC authentication can be configured at the same time on a port. It is assumed that MAC authentication will use an existing MAC address. The following conditions apply for concurrent Web and MAC authentication:

■ A specific MAC address cannot be authenticated by both Web and MAC authentication at the same time.

■ Each new Web/MAC Auth client always initiates a MAC authentication attempt. This same client can also initiate Web authentication at any time before the MAC authentication succeeds. If either authentication succeeds then the other authentication (if in progress) is ended. No further Web/MAC authentication attempts are allowed until the client is deauthenticated.

■ Web and MAC authentications are not allowed on the same port if unauthenticated VLAN (that is, a guest VLAN) is enabled for MAC authentication. An unauthenticated VLAN can't be enabled for MAC authentication if Web and MAC authentication are both enabled on the port.

■ Hitless reauthentication must be of the same type (MAC) that was used for the initial authentication. Non-hitless reauthentication can be of any type.

The remaining Web/MAC functionality, including interactions with 802.1X, remains the same. Web and MAC authentication can be used for different clients on the same port.

Normally, MAC authentication finishes much sooner than Web authentication. However, if Web authentication should complete first, MAC authentication will cease even though it is possible that MAC authentication could succeed. There is no guarantee that MAC authentication ends before Web authentication begins for the client.

Concurrent Web and MAC authentication is backward compatible with all existing user configurations.

Client Options

Web-Auth and MAC-Auth provide a port-based solution in which a port can belong to one, untagged VLAN at a time. The switch allows 32 clients per port.

In the default configuration, the switch blocks access to clients that the RADIUS server does not authenticate. However, you can configure an individual port to provide limited services to unauthorized clients by joining a specified "unauthorized" VLAN during sessions with such clients. The unauthorized VLAN assignment can be the same for all ports, or different, depending on the services and access you plan to allow for unauthenticated clients.

Access to an optional, unauthorized VID is configured in the switch when Web and MAC Authentication are configured on a port.

General Features

Web and MAC Authentication includes the following:

■   On a port configured for Web or MAC Authentication, the switch operates as a port-access authenticator using a RADIUS server and CHAP (Challenge Handshake Authentication Protocol). Inbound traffic is processed by the switch alone, until authentication occurs. Some traffic from the switch is available to an unauthorized client (for example, broadcast or unknown destination packets) before authentication occurs.

■   Proxy servers may not be used by browsers accessing the switch through ports using Web Authentication.

■   You can optionally configure the switch to temporarily assign "authorized" and "unauthorized" VLAN memberships on a per-port basis to provide different services and access to authenticated and unauthenticated clients.

■   Web pages for username and password entry and the display of authorization status are provided when using Web Authentication.

■ You can use the RADIUS server to temporarily assign a port to a static VLAN to support an authenticated client. When a RADIUS server authenticates a client, the switch-port membership during the client's connection is determined according to the following hierarchy:

1. A RADIUS-assigned VLAN

2. An authorized VLAN specified in the Web- or MAC-Auth configuration for the subject port.

3. A static, port-based, untagged VLAN to which the port is configured. A RADIUS-assigned VLAN has priority over switch-port membership in any VLAN.

■ You can allow wireless clients to move between switch ports under Web/MAC Authentication control. Clients may move from one Web authorized port to another or from one MAC authorized port to another. This capability allows wireless clients to move from one access point to another without having to reauthenticate.

■ Unlike 802.1X operation, clients do not need supplicant software for Web or MAC Authentication; only a Web browser (for Web Authentication) or a MAC address (for MAC Authentication).

■ You can use "Show" commands to display session status and port-access configuration settings.

Terminology

Authorized-Client VLAN: Like the Unauthorized-Client VLAN, this is a conventional, static, untagged, port-based VLAN previously configured on the switch by the System Administrator. The intent in using this VLAN is to provide authenticated clients with network access and services. When the client connection terminates, the port drops its membership in this VLAN.

Authentication Server: The entity providing an authentication service to the switch, for example, a RADIUS server.

Authenticator: In HP switch applications, a device that requires a client or device to provide the proper credentials (MAC address, or username and password) before being allowed access to the network.

CHAP: Challenge Handshake Authentication Protocol. Also known as "CHAP-RADIUS".

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

Redirect URL: A System Administrator-specified Web page presented to an authorized client following Web Authentication. HP recommends specifying this URL when configuring Web Authentication on a switch. Refer to **aaa port-access web-based [e] <** *port-list* **> [redirect-url <** *url* **>]** on page 3-22.

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

Unauthorized-Client VLAN: A conventional, static, untagged, port-based VLAN previously configured on the switch by the System Administrator. It is used to provide limited network access and services to clients who are not authenticated.
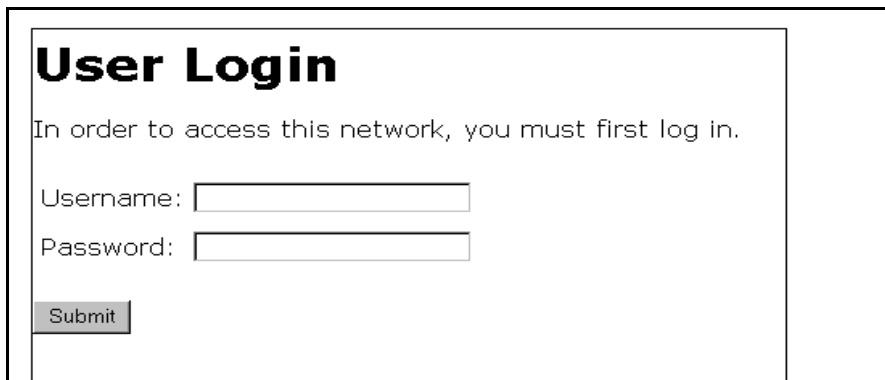
How Web and MAC Authentication Operate

Authenticator Operation

Before gaining access to the network clients first present their authentication credentials to the switch. The switch then verifies the supplied credentials with a RADIUS authentication server. Successfully authenticated clients receive access to the network, as defined by the System Administrator. Clients who fail to authenticate successfully receive no network access or limited network access as defined by the System Administrator.

Web-based Authentication

When a client connects to a Web-Auth enabled port, communication is redirected to the switch. A temporary IP address is assigned by the switch and a login screen is presented for the client to enter their credentials.

User Login

In order to access this network, you must first log in.

Username: [   ]

Password: [   ]

[Submit]

Figure 3-1. Example of User Login Screen

The temporary IP address pool can be specified using the **dhcp-addr** and **dhcp-lease** options of the **aaa port-access web-based** command. If SSL is enabled on the switch and **ssl-login** is enabled on the port the client is redirected to a secure login page (https://...).

The switch passes the supplied username and password to the RADIUS server for authentication.

Authenticating...

Please wait while your credentials are verified.

Figure 3-2. Progress Message During Authentication

If the client is authenticated and the maximum number of clients allowed on the port (**client-limit**) has not been reached, the port is assigned to a static, untagged VLAN for network access. If specified, the client is redirected to a specific URL (**redirect-url**).

Access Granted

You have been authenticated. Please wait while network connection refreshes itself.

Time (sec) Remaining: [20   ]

Figure 3-3. Authentication Completed

The assigned VLAN is determined, in order of priority, as follows:

1. If there is a RADIUS-assigned VLAN, then, for the duration of the client session, the port belongs to this VLAN and temporarily drops all other VLAN memberships.

2. If there is no RADIUS-assigned VLAN, then, for the duration of the client session, the port belongs to the authorized VLAN (**auth-vid** if configured) and temporarily drops all other VLAN memberships.

3. If neither 1 or 2, above, apply, but the port is an untagged member of a statically configured, port-based VLAN, then the port remains in this VLAN.

4. If neither 1, 2, or 3, above, apply, then the client session does not have access to any statically configured, untagged VLANs and client access is blocked.

The assigned port VLAN remains in place until the session ends. Clients may be forced to reauthenticate after a fixed period of time (**reauth-period**) or at any time during a session (**reauthenticate**). An implicit logoff period can be set if there is no activity from the client after a given amount of time (**logoff-period**). In addition, a session ends if the link on the port is lost, requiring reauthentication of all clients. Also, if a client moves from one port to another and client moves have not been enabled (**client-moves**) on the ports, the session ends and the client must reauthenticate for network access. At the end of the session the port returns to its pre-authentication state. Any changes to the port's VLAN memberships made while it is an authorized port take affect at the end of the session.

A client may not be authenticated due to invalid credentials or a RADIUS server timeout. The **max-retries** parameter specifies how many times a client may enter his credentials before authentication fails. The **server-timeout** parameter sets how long the switch waits to receive a response from the RADIUS server before timing out. The **max-requests** parameter specifies how many authentication attempts may result in a RADIUS server timeout before authentication fails. The switch waits a specified amount of time (**quiet-period**) before processing any new authentication requests from the client.

Network administrators may assign unauthenticated clients to a specific static, untagged VLAN (**unauth-vid**), to provide access to specific (guest) network resources. If no VLAN is assigned to unauthenticated clients the port is blocked and no network access is available. Should another client successfully authenticate through that port any unauthenticated clients on the **unauth-vid** are dropped from the port.

MAC Authentication

When a client connects to a MAC-Auth enabled port traffic is blocked.   The switch immediately submits the client's MAC address (in the format specified by the **addr-format**) as its certification credentials to the RADIUS server for authentication.

If the client is authenticated and the maximum number of MAC addresses allowed on the port (**addr-limit**) has not been reached, the port is assigned to a static, untagged VLAN for network access.

The assigned VLAN is determined, in order of priority, as follows:

1. If there is a RADIUS-assigned VLAN, then, for the duration of the client session, the port belongs to this VLAN and temporarily drops all other VLAN memberships.

2. If there is no RADIUS-assigned VLAN, then, for the duration of the client session, the port belongs to the Authorized VLAN (**auth-vid** if configured) and temporarily drops all other VLAN memberships.

3. If neither 1 or 2, above, apply, but the port is an untagged member of a statically configured, port-based VLAN, then the port remains in this VLAN.

4. If neither 1, 2, or 3, above, apply, then the client session does not have access to any statically configured, untagged VLANs and client access is blocked.

The assigned port VLAN remains in place until the session ends. Clients may be forced to reauthenticate after a fixed period of time (**reauth-period**) or at any time during a session (**reauthenticate**). An implicit logoff period can be set if there is no activity from the client after a given amount of time (**logoff-period**). In addition, a session ends if the link on the port is lost, requiring reauthentication of all clients. Also, if a client moves from one port to another and client moves have not been enabled (**addr-moves**) on the ports, the session ends and the client must reauthenticate for network access. At the end of the session the port returns to its pre-authentication state. Any changes to the port's VLAN memberships made while it is an authenticated port take affect at the end of the session.

A client may not be authenticated due to invalid credentials or a RADIUS server timeout. The **server-timeout** parameter sets how long the switch waits to receive a response from the RADIUS server before timing out. The **max-requests** parameter specifies how many authentication attempts may result in

a RADIUS server timeout before authentication fails.   The switch waits a specified amount of time (**quiet-period**) before processing any new authentication requests from the client.

Network administrators may assign unauthenticated clients to a specific static, untagged VLAN (**unauth-vid**), to provide access to specific (guest) network resources. If no VLAN is assigned to unauthenticated clients the port remains in its original VLAN configuration. Should another client successfully authenticate through that port any unauthenticated clients are dropped from the port.

Operating Rules and Notes

■   The switch supports concurrent 802.1X, Web and MAC authentication operation on a port (with up to 32 clients allowed). However, concurrent operation of Web and MAC authentication with other types of authentication on the same port is not supported. That is, the following authentication types are *mutually exclusive* on a given port:

   •   Web and/or MAC Authentication (with or without 802.1X)
   •   MAC lockdown
   •   MAC lockout
   •   Port-Security

■   Order of Precedence for Port Access Management (highest to lowest):

   •   MAC lockout
   •   MAC lockdown or Port Security
   •   Port-based Access Control (802.1X) or Web Authentication or MAC Authentication

Note on Port Access Management

When configuring a port for Web or MAC Authentication, be sure that a higher precedent port access management feature is not enabled on the port. For example, be sure that Port Security is disabled on a port before configuring it for Web or MAC Authentication. If Port Security is enabled on the port this misconfiguration does not allow Web or MAC Authentication to occur.

■ VLANs: If your LAN does not use multiple VLANs, then you do not need to configure VLAN assignments in your RADIUS server or consider using either Authorized or Unauthorized VLANs. If your LAN does use multiple VLANs, then some of the following factors may apply to your use of Web-Auth and MAC-Auth.

• Web-Auth and MAC-Auth operate only with port-based VLANs. Operation with protocol VLANs is not supported, and clients do not have access to protocol VLANs during Web-Auth and MAC-Auth sessions.

• A port can belong to one, untagged VLAN during any client session. Where multiple authenticated clients may simultaneously use the same port, they must all be capable of operating on the same VLAN.

• During an authenticated client session, the following hierarchy determines a port's VLAN membership:
  1. If there is a RADIUS-assigned VLAN, then, for the duration of the client session, the port belongs to this VLAN and temporarily drops all other VLAN memberships.
  2. If there is no RADIUS-assigned VLAN, then, for the duration of the client session, the port belongs to the Authorized VLAN (if configured) and temporarily drops all other VLAN memberships.
  3. If neither 1 or 2, above, apply, but the port is an untagged member of a statically configured, port-based VLAN, then the port remains in this VLAN.
  4. If neither 1, 2, or 3, above, apply, then the client session does not have access to any statically configured, untagged VLANs and client access is blocked.

• After an authorized client session begins on a given port, the port's VLAN membership does not change. If other clients on the same port become authenticated with a different VLAN assignment than the first client, the port blocks access to these other clients until the first client session ends.

• The optional "authorized" VLAN (**auth-vid**) and "unauthorized" VLAN (**unauth-vid**) you can configure for Web- or MAC-based authentication must be statically configured VLANs on the switch. Also, if you configure one or both of these options, any services you want clients in either category to access must be available on those VLANs.

■ Where a given port's configuration includes an unauthorized client VLAN assignment, the port will allow an unauthenticated client session only while there are no requests for an authenticated client session on that port. In this case, if there is a successful request for authentication from an authorized client, the switch terminates the unauthorized-client session and begins the authorized-client session.

■ When a port on the switch is configured for Web or MAC Authentication and is supporting a current session with another device, rebooting the switch invokes a re-authentication of the connection.

■ When a port on the switch is configured as a Web- or MAC-based authenticator, it blocks access to a client that does not provide the proper authentication credentials. If the port configuration includes an optional, unauthorized VLAN (**unauth-vid**), the port is temporarily placed in the unauthorized VLAN if there are no other authorized clients currently using the port with a different VLAN assignment. If an authorized client is using the port with a different VLAN or if there is no unauthorized VLAN configured, the unauthorized client does not receive access to the network.

■ Web- or MAC-based authentication and LACP cannot both be enabled on the same port.

Note on Web/ MAC Authentication and LACP

The switch does not allow Web or MAC Authentication and LACP to both be enabled at the same time on the same port. The switch automatically disables LACP on ports configured for Web or MAC Authentication.

■ Use the **show port-access web-based** commands to display session status, port-access configuration settings, and statistics for Web-Auth sessions.

■ When spanning tree is enabled on a switch that uses 802.1X, Web authentication, or MAC authentication, loops may go undetected. For example, spanning tree packets that are looped back to an edge port will not be processed because they have a different broadcast/multicast MAC address from the client-authenticated MAC address. To ensure that client-authenticated edge ports get blocked when loops occur, you should enable loop protection on those ports. For more information, refer to "Loop Protection" in the chapter titled "Multiple Instance Spanning-Tree Operation" in the *Advanced Traffic Management Guide* for your switch.

General Setup Procedure for Web/MAC Authentication

Do These Steps Before You Configure Web/MAC Authentication

1. Configure a local username and password on the switch for both the Operator (login) and Manager (enable) access levels. (While this is not required for a Web- or MAC-based configuration, HP recommends that you use a local user name and password pair, at least until your other security measures are in place, to protect the switch configuration from unauthorized access.)

2. Determine which ports on the switch you want to operate as authenticators. Note that before you configure Web- or MAC-based authentication on a port operating in an LACP trunk, you must remove the port from the trunk. (refer to the "Note on Web/MAC Authentication and LACP" on page 3-11.)

3. Determine whether any VLAN assignments are needed for authenticated clients.

    a. If you configure the RADIUS server to assign a VLAN for an authenticated client, this assignment overrides any VLAN assignments configured on the switch while the authenticated client session remains active. Note that the VLAN must be statically configured on the switch.

    b. If there is no RADIUS-assigned VLAN, the port can join an "Authorized VLAN" for the duration of the client session, if you choose to configure one. This must be a port-based, statically configured VLAN on the switch.

    c. If there is neither a RADIUS-assigned VLAN or an "Authorized VLAN" for an authenticated client session on a port, then the port's VLAN membership remains unchanged during authenticated client sessions. In this case, configure the port for the VLAN in which you want it to operate during client sessions.

    Note that when configuring a RADIUS server to assign a VLAN, you can use either the VLAN's name or VID. For example, if a VLAN configured in the switch has a VID of 100 and is named **vlan100**, you could configure the RADIUS server to use either "100" or "vlan100" to specify the VLAN.

4. Determine whether to use the optional "Unauthorized VLAN" mode for clients that the RADIUS server does not authenticate. This VLAN must be statically configured on the switch. If you do not configure an "Unauthorized VLAN", the switch simply blocks access to unauthenticated clients trying to use the port.

5. Determine the authentication policy you want on the RADIUS server and configure the server. Refer to the documentation provided with your RADIUS application and include the following in the policy for each client or client device:

   • The CHAP-RADIUS authentication method.

   • An encryption key

   • One of the following:
     – If you are configuring Web-based authentication, include the user name and password for each authorized client.
     – If you are configuring MAC-based authentication, enter the device MAC address in both the username and password fields of the RADIUS policy configuration for that device. Also, if you want to allow a particular device to receive authentication only through a designated port and switch, include this in your policy.

6. Determine the IP address of the RADIUS server(s) you will use to support Web- or MAC-based authentication. (For information on configuring the switch to access RADIUS servers, refer to "Configuring the Switch To Access a RADIUS Server" on page 3-14.)

Configuring the RADIUS Server To Support MAC Authentication

On the RADIUS server, configure the client device authentication in the same way that you would any other client, except:

■ Configure the client device's (hexadecimal) MAC address as both username and password. Be careful to configure the switch to use the same format that the RADIUS server uses. Otherwise, the server will deny access. The switch provides eight format options:

> **aabbccddeeff** (the default format)
>
> **aabbcc-ddeeff**
>
> **aa-bb-cc-dd-ee-ff**
>
> **aa:bb:cc:dd:ee:ff**
>
> **AABBCCDDEEFF**
>
> **AABBCC-DDEEFF**
>
> **AA-BB-CC-DD-EE-FF**

AA:BB:CC:DD:EE:FF

■ If the device is a switch or other VLAN-capable device, use the base MAC address assigned to the device, and not the MAC address assigned to the VLAN through which the device communicates with the authenticator switch. Note that each switch covered by this guide applies a single MAC address to all VLANs configured in the switch. Thus, for a given switch, the MAC address is the same for all VLANs configured on the switch. (Refer to the chapter titled "Static Virtual LANs (VLANs)" in the *Advanced Traffic Management Guide* for your switch.)

Configuring the Switch To Access a RADIUS Server

RADIUS Server Configuration Commands	
radius-server	
[host <*ip-address*>]	below
[key < *global-key-string* >]	below
radius-server host <*ip-address*> key <*server-specific key-string*>	3-15

This section describes the minimal commands for configuring a RADIUS server to support Web-Auth and MAC Auth. For information on other RADIUS command options, refer to chapter 5, "RADIUS Authentication, Authorization 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. You can config- ure up to three RADIUS server addresses. The switch uses the first server it successfully accesses. (Refer to "RADIUS Authentication, Authorization and Account- ing" on page 5-1.)*

[key < *global-key-string* >]

Specifies the global encryption key the switch uses with servers for which the switch does not have a server- specific key assignment (below). This key is optional if all RADIUS server addresses configured in the switch include a server-specific encryption key. (Default: Null.)

Syntax: radius-server host < *ip-address* > key <*server-specific key-string*>
[no] radius-server host < *ip-address* > key

Optional. Specifies an encryption key for use during authentication (or accounting) sessions with the speci- fied 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, above.

*The **no** form of the command removes the key configured for a specific server.*

For example, to configure the switch to access a RADIUS server at IP address 192.168.32.11 using a server-specific shared secret key of '2Pzo22'

```
HP Switch(config)# radius-server host 192.168.32.11 key
2Pzo22
HP Switch(config)# show radius

 Status and Counters - General RADIUS Information

  Deadtime(min) : 0
  Timeout(secs) : 5
  Retransmit Attempts : 3
  Global Encryption Key :

                  Auth  Acct
  Server IP Addr  Port  Port  Encryption Key
  --------------- ----- ----- ----------------------------
```

Figure 3-4. Example of Configuring a Switch To Access a RADIUS Server

Configuring Web Authentication

Overview

1. If you have not already done so, configure a local username and password pair on the switch.

2. Identify or create a redirect URL for use by authenticated clients. HP recommends that you provide a redirect URL when using Web Authentication. If a redirect URL is not specified, Web browser behavior following authentication may not be acceptable.

3. If you plan to use multiple VLANs with Web Authentication, ensure that these VLANs are configured on the switch and that the appropriate port assignments have been made. Also, confirm that the VLAN used by authorized clients can access the redirect URL.

4. Use the **ping** command in the switch console interface to ensure that the switch can communicate with the RADIUS server you have configured to support Web-Auth on the switch.

5. Configure the switch with the correct IP address and encryption key to access the RADIUS server.

6. Configure the switch for Web-Auth:

    a. Configure Web Authentication on the switch ports you want to use.

    b. If the necessary to avoid address conflicts with the secure network, specify the base IP address and mask to be used by the switch for temporary DHCP addresses.The lease length for these temporary IP addresses may also be set.

    c. If you plan to use SSL for logins configure and enable SSL on the switch before you specify it for use with Web-Auth.

    d. Configure the switch to use the redirect URL for authorized clients.

7. Test both authorized and unauthorized access to your system to ensure that Web Authentication works properly on the ports you have configured for port-access using Web Authentication.

Note   Client Web browsers may not use a proxy server to access the network.

Configure the Switch for Web Authentication

Syntax: aaa port-access <*port-list* > controlled-directions <both | in>

> *After you enable web-based authentication on specified ports, you can use the* **aaa port-access controlled-directions** *command to configure how a port transmits traffic before it successfully authenticates a client and enters the authenticated state.*

> **both** *(default)***:** *Incoming and outgoing traffic is blocked on a port configured for web authentication before authentication occurs.*

> **in**: *Incoming traffic is blocked on a port configured for web authentication before authentication occurs. Outgoing traffic with unknown destination addresses is flooded on unauthenticated ports configured for web authentication.*

> ***Prerequisites:*** *As implemented in 802.1X authentication, the disabling of incoming traffic and transmission of outgoing traffic on a web-authenticated egress port in an unauthenticated state (using the* **aaa port-access controlled-directions in** *command) is supported only if:*

> ■   *The 802.1s Multiple Spanning Tree Protocol (MSTP) or 802.1w Rapid Spanning Tree Protocol (RSTP) is enabled on the switch. MSTP and RSTP improve resource utilization while maintaining a loop-free network.*

> *The port is configured as an edge port in the network using the* **spanning-tree edge-port** *command.*

Syntax: aaa port-access <*port-list* > controlled-directions <both | in>
                — Continued —

Notes:

■ *For information on how to configure the prerequisites for using the* **aaa port-access controlled-directions in** *command, see Chapter 4, "Multiple Instance Spanning-Tree Operation" in the* Advanced Traffic Management Guide.

■ *To display the currently configured Controlled Directions value for web-authenticated ports, enter the* **show port-access web-based config** *command.*

■ *The* **aaa port-access controlled-direction in** *command allows Wake-on-LAN traffic to be transmitted on a web-authenticated egress port that has not yet transitioned to the authenticated state; the* **controlled-direction both** *setting prevents Wake-on-LAN traffic to be transmitted on a web-authenticated egress port until authentication occurs.*

   The Wake-on-LAN feature is used by network administrators to remotely power on a sleeping workstation (for example, during early morning hours to perform routine maintenance operations, such as patch management and software updates)

■ *Using the* **aaa port-access controlled-directions in** *command, you can enable the transmission of Wake-on-LAN traffic on unauthenticated egress ports that are configured for any of the following port-based security features:*

   • *802.1X authentication*

   • *MAC authenticatio*n

   • Web authentication

   Because a port can be configured for more than one type of authentication to protect the switch from unauthorized access, the last setting you configure with the **aaa port-access controlled-directions** *command is applied to all authentication methods configured on the switch.*

   For information about how to configure and use 802.1X authentication, refer to Chapter 10, "Configuring Port-Based and Client-Based Access Control (802.1X)".

■ *When a web-authenticated port is configured with the* **controlled-directions in** *setting, eavesdrop prevention is not supported on the port.*

Syntax: aaa port-access web-based dhcp-addr <*ip-address/mask*>

> *Specifies the base address/mask for the temporary IP pool used by DHCP. The base address can be any valid ip address (not a multicast address). Valid mask range value is <255.255.240.0 - 255.255.255.0>. (Default: 192.168.0.0/255.255.255.0)*

Syntax: aaa port-access web-based dhcp-lease <5 - 25>

> *Specifies the lease length, in seconds, of the temporary IP address issued for Web Auth login purposes. (Default: 10 seconds)*

Syntax: [no] aaa port-access web-based [e] < *port-list*>

> *Enables Web-based authentication on the specified ports. Use the* **no** *form of the command to disable Web-based authentication on the specified ports.*

Syntax: aaa port-access web-based [e] < *port-list*> [auth-vid <*vid*>]

aaa port-access web-based [e] < *port-list*> [auth-vid]

> *Specifies the VLAN to use for an authorized client. The Radius server can override the value (accept-response includes a vid). If* **auth-vid** *is* **0**, *no VLAN changes occur unless the RADIUS server supplies one.*
>
> *Use the* **no** *form of the command to set the* **auth-vid** *to* **0**. *(Default: 0).*

Syntax: aaa port-access web-based [e] < *port-list* > [client-limit <1-2>]

> *Specifies the maximum number of authenticated clients to allow on the port. (Default: 1)*

Syntax: [no] aaa port-access web-based [e] < *port-list* > [client-moves]

Allows client moves between the specified ports under Web Auth control. When enabled, the switch allows clients to move without requiring a re-authentication. When disabled, the switch does not allow moves and when one does occur, the user will be forced to re-authenticate. At least two ports (from port(s) and to port(s)) must be specified.

Use the **no** *form of the command to disable client moves between ports under Web Auth control.*
(Default: disabled – no moves allowed)

Syntax: aaa port-access web-based [e] < *port-list* > [logoff-period]
<60-9999999>]

Specifies the period, in seconds, that the switch enforces for an implicit logoff. This parameter is equivalent to the MAC age interval in a traditional switch sense. If the switch does not see activity after a logoff-period interval, the client is returned to its pre-authentication state. (Default: 300 seconds)

Syntax: aaa port-access web-based [e] < *port-list* > [max-requests <1-10>]

Specifies the number of authentication attempts that must time-out before authentication fails.
(Default: 2)

Syntax: aaa port-access web-based [e] < *port-list* > [max-retries <1-10>]

Specifies the number of the number of times a client can enter their user name and password before authentication fails. This allows the reentry of the user name and password if necessary.
(Default: 3)

Syntax: aaa port-access web-based [e] < *port-list* > [quiet-period <1 - 65535>]

> *Specifies the time period, in seconds, the switch should wait before attempting an authentication request for a client that failed authentication.*
> *(Default: 60 seconds)*

Syntax: aaa port-access web-based [e] < *port-list* > [reauth-period <0 - 9999999>]

> *Specifies the time period, in seconds, the switch enforces on a client to re-authenticate. When set to* **0**, *reauthentication is disabled. (Default: 300 seconds)*

Syntax: aaa port-access web-based [e] < *port-list* > [reauthenticate]

> *Forces a reauthentication of all attached clients on the port.*

Syntax: aaa port-access web-based [e] < *port-list* > [redirect-url <*url*>]

         no aaa port-access web-based [e] < *port-list* > [redirect-url]

> *Specifies the URL that a user is redirected to after a successful login. Any valid, fully-formed URL may be used, for example, http://welcome-server/welcome.htm or http://192.22.17.5. HP recommends that you provide a redirect URL when using Web Authentication.*

> *Use the* **no** *form of the command to remove a specified redirect URL.*
> *(Default: There is no default URL. Browser behavior for authenticated clients may not be acceptable.)*

Syntax: aaa port-access web-based [e] < *port-list* > [server-timeout <1 - 300>]

> *Specifies the period, in seconds, the switch waits for a server response to an authentication request. Depending on the current* **max-requests** *value, the switch sends a new attempt or ends the authentication session. (Default: 30 seconds)*

Syntax: [no] aaa port-access web-based [e] < *port-list* > [ssl-login]

>*Enables or disables SSL login (https on port 443). SSL must be enabled on the switch.*
>
>*If SSL login is* enabled, *a user is redirected to a secure page, where they enter their username and password. If SSL login is* disabled, *a user is not redirected to a secure page to enter their credentials.*
>
>*Use the* **no** *form of the command to disable SSL login. (Default: disabled)*

Syntax: aaa port-access web-based [e] < *port-list* > [unauth-vid <*vid*>]

   no aaa port-access web-based [e] < *port-list* > [unauth-vid]

>*Specifies the VLAN to use for a client that fails authentication. If* **unauth-vid** *is* **0**, *no VLAN changes occur.*
>
>*Use the* **no** *form of the command to set the* **unauth-vid** *to* **0**. *(Default: 0)*

Configuring MAC Authentication on the Switch

Overview

1. If you have not already done so, configure a local username and password pair on the switch.

2. If you plan to use multiple VLANs with MAC Authentication, ensure that these VLANs are configured on the switch and that the appropriate port assignments have been made.

3. Use the **ping** command in the switch console interface to ensure that the switch can communicate with the RADIUS server you have configured to support MAC-Auth on the switch.

4. Configure the switch with the correct IP address and encryption key to access the RADIUS server.

5. Configure the switch for MAC-Auth:

    a. Configure MAC Authentication on the switch ports you want to use.

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

Configure the Switch for MAC-Based Authentication

Command	Page
Configuration Level	
aaa port-access mac-based addr-format	3-24
[no] aaa port-access mac-based [e] < *port-list* >	3-25
[addr-limit]	3-25
[addr-moves]	3-25
[auth-vid]	3-25
[logoff-period]	3-26
[max-requests]	3-26
[quiet-period]	3-26
[reauth-period]	3-26
[reauthenticate]	3-26
[server-timeout]	3-26
[unauth-vid]	3-27

Syntax: aaa port-access mac-based addr-format
             <no-delimiter|single-dash|multi-dash|multi-colon>

> *Specifies the MAC address format to be used in the RADIUS request message. This format must match the format used to store the MAC addresses in the RADIUS server. (Default: no-delimiter)*
>
> **no-delimiter** — *specifies an aabbccddeeff format.*
>
> **single-dash** — *specifies an aabbcc-ddeeff format.*
>
> **multi-dash** — *specifies an aa-bb-cc-dd-ee-ff format.*
>
> **multi-colon** — *specifies an aa:bb:cc:dd:ee:ff format.*
>
> **no-delimiter-uppercase** — *specifies an AABBCCDDEEFF format.*
>
> **single-dash-uppercase** — *specifies an AABBCC-DDEEFF format*

multi-dash-uppercase — *specifies an AA-BB-CC-DD-EE-FF format*

multi-colon-uppercase — *specifies an AA:BB:CC:DD:EE:FF format.*

Syntax: [no] aaa port-access mac-based [e] < *port-list* >

> *Enables MAC-based authentication on the specified ports. Use the* **no** *form of the command to disable MAC-based authentication on the specified ports.*

Syntax: aaa port-access mac-based [e] < *port-list* > [addr-limit <1-2>]

> *Specifies the maximum number of authenticated MACs to allow on the port. (Default: 1)*

Syntax: [no] aaa port-access mac-based [e] < *port-list* > [addr-moves]

> *Allows client moves between the specified ports under MAC Auth control. When enabled, the switch allows addresses to move without requiring a re-authentication. When disabled, the switch does not allow moves and when one does occur, the user will be forced to re-authenticate. At least two ports (from port(s) and to port(s)) must be specified.*
>
> *Use the* **no** *form of the command to disable MAC address moves between ports under MAC Auth control. (Default: disabled – no moves allowed)*

Syntax: aaa port-access mac-based [e] < *port-list* > [auth-vid <*vid*>]

no aaa port-access mac-based [e] < *port-list* > [auth-vid]

> *Specifies the VLAN to use for an authorized client. The Radius server can override the value (accept-response includes a vid). If* **auth-vid** *is* **0***, no VLAN changes occur unless the RADIUS server supplies one.*
>
> *Use the* **no** *form of the command to set the* **auth-vid** *to* **0***. (Default: 0).*

Syntax: aaa port-access mac-based [e] < *port-list* >
                [logoff-period] <60-9999999>]

> *Specifies the period, in seconds, that the switch
> enforces for an implicit logoff. This parameter is
> equivalent to the MAC age interval in a traditional
> switch sense. If the switch does not see activity after a
> logoff-period interval, the client is returned to its pre-
> authentication state. (Default: 300 seconds)*

Syntax: aaa port-access mac-based [e] < *port-list* > [max-requests <1-10>]

> *Specifies the number of authentication attempts that
> must time-out before authentication fails.
> (Default: 2)*

Syntax: aaa port-access mac-based [e] < *port-list* > [quiet-period <1 - 65535>]

> *Specifies the time period, in seconds, the switch should
> wait before attempting an authentication request for
> a MAC address that failed authentication.
> (Default: 60 seconds)*

Syntax: aaa port-access mac-based [e] < *port-list* > [reauth-period <0 - 9999999>]

> *Specifies the time period, in seconds, the switch
> enforces on a client to re-authenticate. When set to **0**,
> reauthentication is disabled. (Default: 300 seconds)*

Syntax: aaa port-access mac-based [e] < *port-list* > [reauthenticate]

> *Forces a reauthentication of all attached clients on the
> port.*

Syntax: aaa port-access mac-based [e] < *port-list* > [server-timeout <1 - 300>]

> *Specifies the period, in seconds, the switch waits for a
> server response to an authentication request. Depend-
> ing on the current **max-requests** value, the switch sends
> a new attempt or ends the authentication session.
> (Default: 30seconds)*

Syntax: aaa port-access mac-based [e] < *port-list* > [unauth-vid <*vid*>]

no aaa port-access mac-based [e] < *port-list* > [unauth-vid]

> *Specifies the VLAN to use for a client that fails authen-tication. If **unauth-vid** is **0**, no VLAN changes occur.*

> *Use the **no** form of the command to set the **unauth-vid** to **0**. (Default: 0)*

Show Status and Configuration of Web-Based Authentication

Command	Page
show port-access [*port-list*] web-based	3-27
[clients]	3-27
[config]	3-28
[config [auth-server]]	3-28
[config [web-server]]	3-28
show port-access *port-list* web-based config detail	3-28

Syntax: show port-access [*port-list*] web-based

> *Shows the status of all Web-Authentication enabled ports or the specified ports. The number of authorized and unauthorized clients is listed for each port, as well as its current VLAN ID. Ports without Web Authenti-cation enabled are not listed.*

Syntax: show port-access [*port-list*] web-based [clients]

> *Shows the port address, Web address, session status, and elapsed session time for attached clients on all ports or the specified ports. Ports with multiple clients have an entry for each attached client. Ports without any attached clients are not listed.*

Syntax: show port-access [*port-list*] web-based [config]

> *Shows Web Authentication settings for all ports or the specified ports, including the temporary DHCP base address and mask. The authorized and unauthorized VLAN IDs are shown. If the authorized or unauthorized VLAN ID is 0 then no VLAN change is made, unless the RADIUS server supplies one.*

Syntax: show port-access [*port-list*] web-based [config [auth-server]]

> *Shows Web Authentication settings for all ports or the specified ports, along with the RADIUS server specific settings for the timeout wait, the number of timeout failures before authentication fails, and the length of time between authentication requests.*

Syntax: show port-access [*port-list*] web-based [config [web-server]]

> *Shows Web Authentication settings for all ports or the specified ports, along with the Web specific settings for password retries, SSL login status, and a redirect URL, if specified.*

Syntax: show port-access *port-list* web-based config detail

> *Shows all Web Authentication settings, including the Radius server specific settings for the specified ports.*

Show Status and Configuration of MAC-Based Authentication

Command	Page
show port-access [*port-list*] mac-based	3-29
[clients]	3-29
[config]	3-29
[config [auth-server]]	3-29
show port-access *port-list* mac-based config detail	3-29

Syntax: show port-access [*port-list*] mac-based

> *Shows the status of all MAC-Authentication enabled ports or the specified ports. The number of authorized and unauthorized clients is listed for each port, as well as its current VLAN ID. Ports without MAC Authentication enabled are not listed.*

Syntax: show port-access [*port-list*] mac-based [clients]

> *Shows the port address, MAC address, session status, and elapsed session time for attached clients on all ports or the specified ports. Ports with multiple clients have an entry for each attached client. Ports without any attached clients are not listed.*

Syntax: show port-access [*port-list*] mac-based [config]

> *Shows MAC Authentication settings for all ports or the specified ports, including the MAC address format being used. The authorized and unauthorized VLAN IDs are shown. If the authorized or unauthorized VLAN ID is 0 then no VLAN change is made, unless the RADIUS server supplies one.*

Syntax: show port-access [*port-list*] mac-based [config [auth-server]]

> *Shows MAC Authentication settings for all ports or the specified ports, along with the Radius server specific settings for the timeout wait, the number of timeout failures before authentication fails, and the length of time between authentication requests.*

Syntax: show port-access *port-list* mac-based config detail

> *Shows all MAC Authentication settings, including the Radius server specific settings for the specified ports.*

Show Client Status

The table below shows the possible client status information that may be reported by a Web-based or MAC-based '**show... clients'** command.

Reported Status	Available Network Connection	Possible Explanations
authenticated	Authorized VLAN	Client authenticated. Remains connected until logoff-period or reauth-period expires.
authenticating	Switch only	Pending RADIUS request.
rejected-no vlan	No network access	1. Invalid credentials supplied. 2. RADIUS Server difficulties. See log file. 3. If unauth-vid is specified it cannot be successfully applied to the port. An authorized client on the port has precedence.
rejected-unauth vlan	Unauthorized VLAN only	1. Invalid credentials supplied. 2. RADIUS Server difficulties. See log file.
timed out-no vlan	No network access	RADIUS request timed out. If unauth-vid is specified it cannot be successfully applied to the port. An authorized client on the port has precedence. Credentials resubmitted after quiet-period expires.
timed out-unauth vlan	Unauthorized VLAN only	RADIUS request timed out. After the quiet-period expires credentials are resubmitted when client generates traffic.
unauthenticated	Switch only	Waiting for user credentials.

TACACS+ Authentication

Overview

Feature	Default	Menu	CLI	Web
view the switch's authentication configuration	n/a	—	page 4-7	—
view the switch's TACACS+ server contact configuration	n/a	—	page 4-7	—
configure the switch's authentication methods	disabled	—	page 4-8	—
configure the switch to contact TACACS+ server(s)	disabled	—	page 4-13	—

TACACS+ authentication enables you to use a central server to allow or deny access to the switches covered in this guide (and other TACACS-aware devices) in your network. This means that you can use a central database to create multiple unique username/password sets with associated privilege levels for use by individuals who have reason to access the switch from either the switch's console port (local access) or Telnet (remote access).
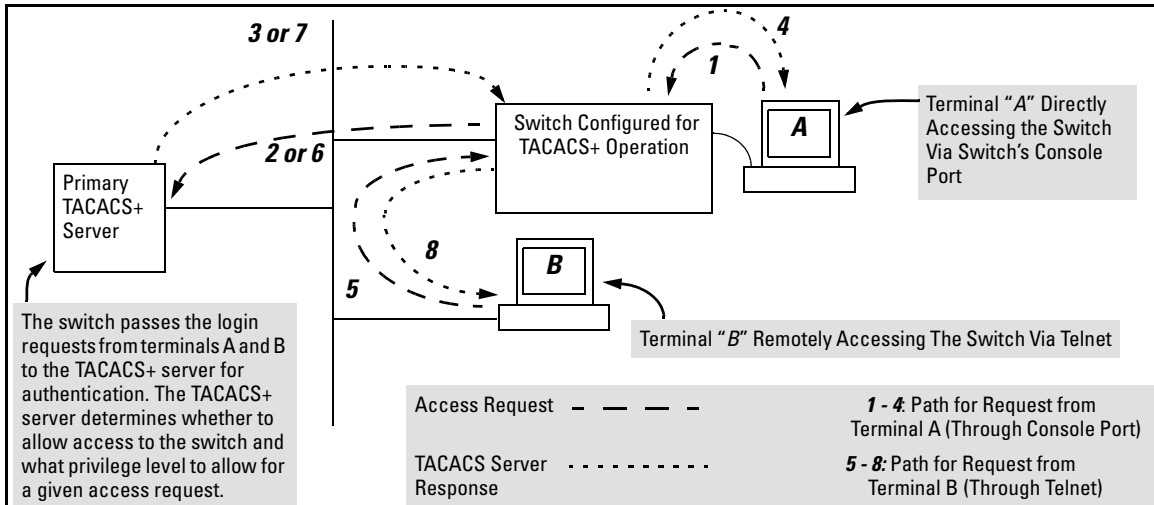
Figure 4-1. Example of TACACS+ Operation

TACACS+ in the switches covered in this guide manages authentication of logon attempts through either the Console port or Telnet. TACACS+ uses an authentication hierarchy consisting of (1) remote passwords assigned in a TACACS+ server and (2) local passwords configured on the switch. That is, with TACACS+ configured, the switch first tries to contact a designated TACACS+ server for authentication services. If the switch fails to connect to

any TACACS+ server, it defaults to its own locally assigned passwords for authentication control if it has been configured to do so. For both Console and Telnet access you can configure a login (read-only) and an enable (read/write) privilege level access.

TACACS+ does not affect web browser interface access. See "Controlling Web Browser Interface Access" on page 4-22.

Terminology Used in TACACS Applications:

- **NAS (Network Access Server):** This is an industry term for a TACACS-aware device that communicates with a TACACS server for authentication services. Some other terms you may see in literature describing TACACS operation are *communication server*, *remote access server*, or *terminal server*. These terms apply to a switch when TACACS+ is enabled on the switch (that is, when the switch is TACACS-aware).

- **TACACS+ Server:** The server or management station configured as an access control server for TACACS-enabled devices. To use TACACS+ with a switch covered in this guide and any other TACACS-capable devices in your network, you must purchase, install, and configure a TACACS+ server application on a networked server or management station in the network. The TACACS+ server application you install will provide various options for access control and access notifications. For more on the TACACS+ services available to you, see the documentation provided with the TACACS+ server application you will use.

- **Authentication:** The process for granting user access to a device through entry of a user name and password and comparison of this username/password pair with previously stored username/password data. Authentication also grants levels of access, depending on the privileges assigned to a user name and password pair by a system administrator.

  - **Local Authentication:** This method uses username/password pairs configured locally on the switch; one pair each for manager-level and operator-level access to the switch. You can assign local usernames and passwords through the CLI or web browser interface. (Using the menu interface you can assign a local password, but not a username.) Because this method assigns passwords to the switch instead of to individuals who access the switch, you must distribute the password information on each switch to everyone who needs to access the switch, and you must configure

and manage password protection on a per-switch basis. (For more on local authentication, refer to chapter 2, "Configuring Local Password Security".)

- **TACACS+ Authentication:** This method enables you to use a TACACS+ server in your network to assign a unique password, user name, and privilege level to each individual or group who needs access to one or more switches or other TACACS-aware devices. This allows you to administer primary authentication from a central server, and to do so with more options than you have when using only local authentication. (You will still need to use local authentication as a backup if your TACACS+ servers become unavailable.) This means, for example, that you can use a central TACACS+ server to grant, change, or deny access to a specific individual on a specific switch instead of having to change local user name and password assignments on the switch itself, and then have to notify other users of the change.

General System Requirements

To use TACACS+ authentication, you need the following:

- A TACACS+ server application installed and configured on one or more servers or management stations in your network. (There are several TACACS+ software packages available.)
- A switch configured for TACACS+ authentication, with access to one or more TACACS+ servers.

Notes

The effectiveness of TACACS+ security depends on correctly using your TACACS+ server application. For this reason, HP recommends that you thoroughly test all TACACS+ configurations used in your network.

TACACS-aware HP switches include the capability of configuring multiple backup TACACS+ servers. HP recommends that you use a TACACS+ server application that supports a redundant backup installation. This allows you to configure the switch to use a backup TACACS+ server if it loses access to the first-choice TACACS+ server.

TACACS+ does not affect web browser interface access. Refer to "Controlling Web Browser Interface Access When Using TACACS+ Authentication" on page 4-22.

General Authentication Setup Procedure

It is important to test the TACACS+ service before fully implementing it. Depending on the process and parameter settings you use to set up and test TACACS+ authentication in your network, you could accidentally lock all users, including yourself, out of access to a switch. While recovery is simple, it may pose an inconvenience that can be avoided. To prevent an unintentional lockout on the switch, use a procedure that configures and tests TACACS+ protection for one access type (for example, Telnet access), while keeping the other access type (console, in this case) open in case the Telnet access fails due to a configuration problem. The following procedure outlines a general setup procedure.

Note

If a complete access lockout occurs on the switch as a result of a TACACS+ configuration, see "Troubleshooting TACACS+ Operation" in the Trouble-shooting chapter of the *Management and Configuration Guide* for your switch.

1. Familiarize yourself with the requirements for configuring your TACACS+ server application to respond to requests from the switch. (Refer to the documentation provided with the TACACS+ server soft-ware.) This includes knowing whether you need to configure an encryp-tion key. (See "Using the Encryption Key" on page 4-20.)

2. Determine the following:

   - The IP address(es) of the TACACS+ server(s) you want the switch to use for authentication. If you will use more than one server, determine which server is your first-choice for authentication services.
   - The encryption key, if any, for allowing the switch to communicate with the server. You can use either a global key or a server-specific key, depending on the encryption configuration in the TACACS+ server(s).
   - The number of log-in attempts you will allow before closing a log-in session. (Default: 3)
   - The period you want the switch to wait for a reply to an authentication request before trying another server.
   - The username/password pairs you want the TACACS+ server to use for controlling access to the switch.
   - The privilege level you want for each username/password pair administered by the TACACS+ server for controlling access to the switch.
   - The username/password pairs you want to use for local authentication (one pair each for Operator and Manager levels).

3. Plan and enter the TACACS+ server configuration needed to support TACACS+ operation for Telnet access (login and enable) to the switch. This includes the username/password sets for logging in at the Operator (read-only) privilege level and the sets for logging in at the Manager (read/write) privilege level.

Note on Privilege Levels   When a TACACS+ server authenticates an access request from a switch, it includes a privilege level code for the switch to use in determining which privilege level to grant to the terminal requesting access. The switch interprets a privilege level code of "15" as authorization for the Manager (read/write) privilege level access. Privilege level codes of 14 and lower result in Operator (read-only) access. Thus, when configuring the TACACS+ server response to a request that includes a username/password pair that should have Manager privileges, you must use a privilege level of 15. For more on this topic, refer to the documentation you received with your TACACS+ server application.

If you are a first-time user of the TACACS+ service, HP recommends that you configure only the minimum feature set required by the TACACS+ application to provide service in your network environment. After you have success with the minimum feature set, you may then want to try additional features that the application offers.

4. Ensure that the switch has the correct local username and password for Manager access. (If the switch cannot find any designated TACACS+ servers, the local manager and operator username/password pairs are always used as the secondary access control method.)

Caution   You should ensure that the switch has a local Manager password. Otherwise, if authentication through a TACACS+ server fails for any reason, then unauthorized access will be available through the console port or Telnet.

5. Using a terminal device connected to the switch's console port, configure the switch for TACACS+ authentication *only* for **telnet login** access and **telnet enable** access. At this stage, do not configure TACACS+ authentication for console access to the switch, as you may need to use the console for access if the configuration for the Telnet method needs debugging.

6. Ensure that the switch is configured to operate on your network and can communicate with your first-choice TACACS+ server. (At a minimum, this requires IP addressing and a successful **ping** test from the switch to the server.)

7. On a remote terminal device, use Telnet to attempt to access the switch. If the attempt fails, use the console access to check the TACACS+ configuration on the switch. If you make changes in the switch configu-

ration, check Telnet access again. If Telnet access still fails, check the configuration in your TACACS+ server application for mis-configurations or missing data that could affect the server's interoperation with the switch.

8. After your testing shows that Telnet access using the TACACS+ server is working properly, configure your TACACS+ server application for console access. Then test the console access. If access problems occur, check for and correct any problems in the switch configuration, and then test console access again. If problems persist, check your TACACS+ server application for mis-configurations or missing data that could affect the console access.

9. When you are confident that TACACS+ access through both Telnet and the switch's console operates properly, use the **write memory** command to save the switch's running-config file to flash.

Configuring TACACS+ on the Switch

Before You Begin

If you are new to TACACS+ authentication, HP recommends that you read the "General Authentication Setup Procedure" on page 4-4 and configure your TACACS+ server(s) before configuring authentication on the switch.

The switch offers three command areas for TACACS+ operation:

■ **show authentication** and **show tacacs:** Displays the switch's TACACS+ configuration and status.

■ **aaa authentication:** A command for configuring the switch's authentication methods

■ **tacacs-server:** A command for configuring the switch's contact with TACACS+ servers

CLI Commands Described in this Section

Command	Page
show authentication	4-7
show tacacs	4-7
aaa authentication	4-8 through 4-13

Viewing the Switch's Current Authentication Configuration

This command lists the number of login attempts the switch allows in a single login session, and the primary/secondary access methods configured for each type of access.

Syntax: show authentication

This example shows the default authentication configuration.

```
Switch> show authentication
 Status and Counters - Authentication Information
  Login Attempts : 3

              Login      Login      Enable     Enable
  Access Task Primary    Secondary  Primary    Secondary
  ---------- -------- ---------- -------- ---------
  Console     Local      None       Local      None
  Telnet      Local      None       Local      None
```

Configuration for login and enable access to the switch through the switch console port.

Configuration for login and enable access to the switch through Telnet.

Figure 4-2. Example Listing of the Switch's Authentication Configuration

Viewing the Switch's Current TACACS+ Server Contact Configuration

This command lists the timeout period, encryption key, and the IP addresses of the first-choice and backup TACACS+ servers the switch can contact.

Syntax: show tacacs

For example, if the switch was configured for a first-choice and two backup TACACS+ server addresses, the default timeout period, and **paris-1** for a (global) encryption key, **show tacacs** would produce a listing similar to the following:

```
                    HP Switch# show tacacs
                     Status and Counters - TACACS Information
First-Choice          Timeout : 5
TACACS+ Server        Encryption Key : paris-1
                      Server IP Addr   Opens  Closes Ahorts Errors Pkts Rx Pkts Tx
Second-Choice         --------------  ------ ------ ------ ------ ------- -------
TACACS+ Server        10.30.248 .100   0       0      0      0      0       0
                      10.30.248 .156   0       0      0      0      0       0
Third-Choice          10.30.248 .105   0       0      0      0      0       0
TACACS+ Server
```
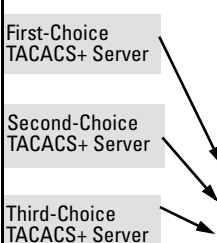
Figure 4-3. Example of the Switch's TACACS+ Configuration Listing

Configuring the Switch's Authentication Methods

The **aaa authentication** command configures access control for the following access methods:

■   Console

■   Telnet

■   SSH

■   Web

■   Port-access (802.1X)

However, TACACS+ authentication is only used with the console, Telnet, or SSH access methods. The command specifies whether to use a TACACS+ server or the switch's local authentication, or (for some secondary scenarios) no authentication (meaning that if the primary method fails, authentication is denied). The command also reconfigures the number of access attempts to allow in a session if the first attempt uses an incorrect username/password pair.

Using the Privilege-Mode Option for Login

When using TACACS+ to control user access to the switch, you must first login with your username at the Operator privilege level using the password for Operator privileges, and then login again with the same username but using the Manger password to obtain Manager privileges. You can avoid this double login process by entering the **privilege-mode** option with the **aaa authentication**

login command to enable TACACS+ for a single login. The switch authenti-
cates your username/password, then requests the privilege level (Operator or
Manager) that was configured on the TACACS+ server for this username/
password. The TACACS+ server returns the allowed privilege level to the
switch. You are placed directly into Operator or Manager mode, depending on
your privilege level.

```
HP Switch(config) aaa authentication login privilege-mode
```

The **no** version of the command disables TACACS+ single login capability.

Syntax: aaa authentication

< console | telnet | ssh | web | port-access >

Selects the access method for configuration.

< enable>

*The server grants privileges at the Manager privilege
level.*

<login [privilege-mode] >

*The server grants privileges at the Operator privilege
level. If the* **privilege-mode** *option is entered, TACACS+ is
enabled for a single login. The authorized privilege level
(Operator or Manager) is returned to the switch by the
TACACS+ server.*
Default: Single login disabled.

< local | tacacs | radius >

Selects the type of security access:

local — *Authenticates with the Manager and Operator
password you configure in the switch.*
tacacs — *Authenticates with a password and other
data configured on a TACACS+ server.*
radius — *Authenticates with a password and other
data configured on a RADIUS server.*

[< local | none >]

*If the primary authentication method fails, determines
whether to use the local password as a secondary method
or to disallow access.*

aaa authentication num-attempts < 1-10 >

*Specifies the maximum number of login attempts allowed in
the current session. Default: 3*

Authentication Parameters

Table 4-1. AAA Authentication Parameters

Name	Default	Range	Function
console, Telnet, SSH, web or port-access	n/a	n/a	Specifies the access method used when authenticating. TACACS+ authentication only uses the console, Telnet or SSH access methods.
enable	n/a	n/a	Specifies the Manager (read/write) privilege level for the access method being configured.
login <privilege-mode>	privilege-mode disabled	n/a	**login:** Specifies the Operator (read-only) privilege level for the access method being configured. The **privilege-mode** option enables TACACS+ for a single login. The authorized privilege level (Operator or Manager) is returned to the switch by the TACACS+ server.
local *- or -* tacacs	local	n/a	Specifies the primary method of authentication for the access method being configured. **local:** Use the username/password pair configured locally in the switch for the privilege level being configured **tacacs:** Use a TACACS+ server.
local *- or -* none	none	n/a	Specifies the secondary (backup) type of authentication being configured. **local:** The username/password pair configured locally in the switch for the privilege level being configured **none:** No secondary type of authentication for the specified method/privilege path. **(Available only if the primary method of authentication for the access being configured is local.)** **Note:** If you do not specify this parameter in the command line, the switch automatically assigns the secondary method as follows: • If the primary method is **tacacs**, the **only** secondary method is **local**. • If the primary method is **local**, the default secondary method is **none**.
num-attempts	3	1 - 10	In a given session, specifies how many tries at entering the correct username/password pair are allowed before access is denied and the session terminated.

Configuring the TACACS+ Server for Single Login

In order for the single login feature to work correctly, you need to check some entries in the User Setup on the TACACS+ server.

In the User Setup, scroll to the Advanced TACACS+ Settings section. Make sure the radio button for "Max Privilege for any AAA Client" is checked and the level is set to 15, as shown in Figure 4-4. Privileges are represented by the numbers 0 through 15, with zero allowing only Operator privileges (and requiring two logins) and 15 representing root privileges. The root privilege level is the only level that will allow Manager level access on the switch.
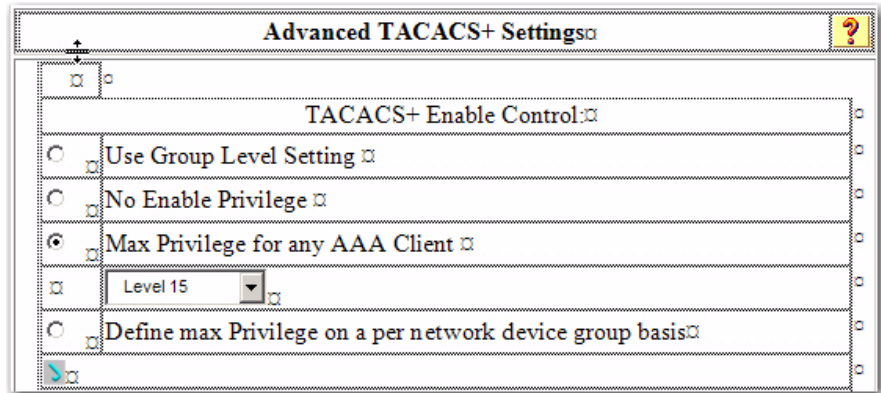
Figure 4-4. Advanced TACACS+ Settings Section of the TACACS+ Server User Setup

Then scroll down to the section that begins with "Shell" (See Figure 4-5). Check the Shell box.

Check the Privilege level box and set the privilege level to 15 to allow "root" privileges. This allows you to use the single login option.
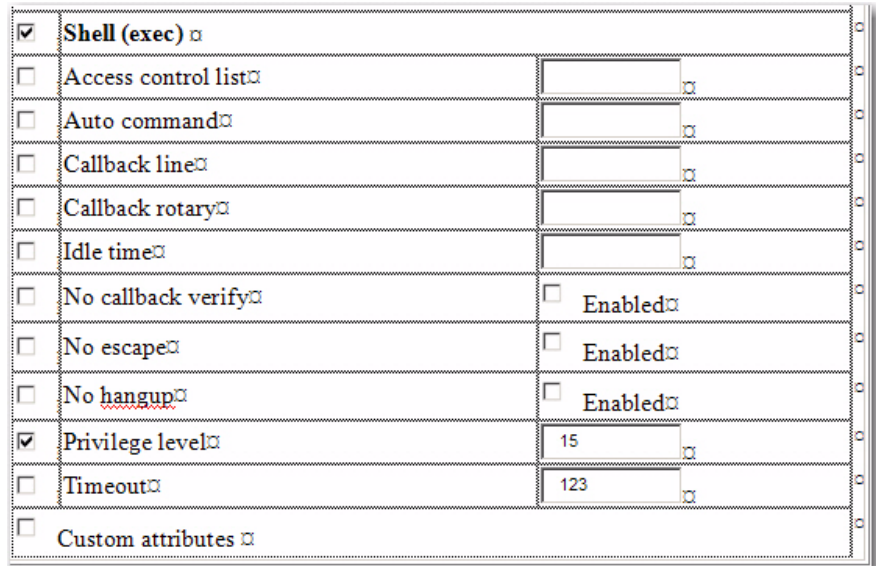
Figure 4-5. The Shell Section of the TACACS+ Server User Setup

As shown in the next table, login and enable access is always available locally through a direct terminal connection to the switch's console port. However,

for Telnet access, you can configure TACACS+ to deny access if a TACACS+ server goes down or otherwise becomes unavailable to the switch.

Table 4-2. Primary/Secondary Authentication Table

Access Method and Privilege Level	Authentication Options		Effect on Access Attempts
	Primary	**Secondary**	
Console — Login	local	none*	Local username/password access only.
	tacacs	local	If Tacacs+ server unavailable, uses local username/password access.
Console — Enable	local	none	Local username/password access only.
	tacacs	local	If Tacacs+ server unavailable, uses local username/password access.
Telnet — Login	local	none*	Local username/password access only.
	tacacs	local	If Tacacs+ server unavailable, uses local username/password access.
	tacacs	none	If Tacacs+ server unavailable, denies access.
Telnet — Enable	local	none	Local username/password access only.
	tacacs	local	If Tacacs+ server unavailable, uses local username/password access.
	tacacs	none	If Tacacs+ server unavailable, denies access.

Caution Regarding the Use of Local for Login Primary Access

During local authentication (which uses passwords configured in the switch instead of in a TACACS+ server), the switch grants read-only access if you enter the Operator password, and read-write access if you enter the Manager password. For example, if you configure authentication on the switch with Telnet Login Primary as Local and Telnet Enable Primary as Tacacs, when you attempt to Telnet to the switch, you will be prompted for a local password. If you enter the switch's local Manager password (or, if there is no local Manager password configured in the switch) you can bypass the TACACS+ server authentication for Telnet Enable Primary and go directly to read-write (Manager) access. Thus, for either the Telnet or console access method, configuring Login Primary for Local authentication while configuring Enable Primary for TACACS+ authentication is not recommended, as it defeats the purpose of using the TACACS+ authentication. If you want Enable Primary log-in attempts to go to a TACACS+ server, then you should configure both Login Primary and Enable Primary for Tacacs authentication instead of configuring Login Primary to Local authentication.

For example, here is a set of access options and the corresponding commands to configure them:

Console Login (Operator or Read-Only) Access: Primary using TACACS+ server.
   Secondary using Local.

```
HP Switch(config)# aaa authentication console login tacacs local
```

Console Enable (Manager or Read/Write) Access: Primary using TACACS+ server.
   Secondary using Local.

```
HP Switch(config)# aaa authentication console enable tacacs local
```

Telnet Login (Operator or Read-Only) Access: Primary using TACACS+ server.
   Secondary using Local.

```
HP Switch(config)# aaa authentication Telnet login tacacs local
```

Telnet Enable (Manager or Read/Write Access: Primary using TACACS+ server.
   Secondary using Local.

```
HP Switch(config)# aaa authentication telnet enable tacacs local
```

Deny Access and Close the Session After Failure of Two Consecutive Username/Password Pairs:

```
HP Switch(config)# aaa authentication num-attempts 2
```

Configuring the Switch's TACACS+ Server Access

The tacacs-server command configures these parameters:

- **The host IP address(es)** for up to three TACACS+ servers; one first-choice and up to two backups. Designating backup servers provides for a continuation of authentication services in case the switch is unable to contact the first-choice server.

- **An optional encryption key.** This key helps to improve security, and must match the encryption key used in your TACACS+ server application. In some applications, the term "secret key" or "secret" may be used instead of "encryption key". If you need only one encryption key for the switch to use in all attempts to authenticate through a TACACS+ server, configure a global key. However, if the switch is configured to access multiple TACACS+ servers having different encryption keys, you can configure the switch to use different encryption keys for different TACACS+ servers.

- **The timeout value** in seconds for attempts to contact a TACACS+ server. If the switch sends an authentication request, but does not receive a response within the period specified by the timeout value, the switch resends the request to the next server in its Server IP Addr list, if any. If

the switch still fails to receive a response from any TACACS+ server, it reverts to whatever secondary authentication method was configured using the **aaa authentication** command (local or none; see "Configuring the Switch's Authentication Methods" on page 4-8.)

Note

As described under "General Authentication Setup Procedure" on page 4-4, HP recommends that you configure, test, and troubleshoot authentication via Telnet access before you configure authentication via console port access. This helps to prevent accidentally locking yourself out of switch access due to errors or problems in setting up authentication in either the switch or your TACACS+ server.

Syntax: tacacs-server host < *ip-addr* > [key < *key-string* >]

Adds a TACACS+ server and optionally assigns a server-specific encryption key.

[no] tacacs-server host < *ip-addr* >

Removes a TACACS+ server assignment (including its server-specific encryption key, if any).

tacacs-server key <*key-string*>

Enters the optional global encryption key.

[no] tacacs-server key

Removes the optional global encryption key. (Does not affect any server-specific encryption key assignments.)

tacacs-server timeout < 1-255 >

Changes the wait period for a TACACS server response. (Default: 5 seconds.)

Note

Encryption keys configured in the switch must exactly match encryption keys configured in TACACS+ servers the switch will attempt to use for authentication.

If you configure a global encryption key, the switch uses it only with servers for which you have not also configured a server-specific key. Thus, a global key is more useful where the TACACS+ servers you are using all have an identical key, and server-specific keys are necessary where different TACACS+ servers have different keys.

If TACACS+ server "X" does not have an encryption key assigned for the switch, then configuring either a global encryption key or a server-specific key in the switch for server "X" will block authentication support from server "X".

Name	Default	Range
host *<ip-addr>* [key *<key-string>*	none	n/a

Specifies the IP address of a device running a TACACS+ server application. Optionally, can also specify the unique, per-server encryption key to use when each assigned server has its own, unique key. For more on the encryption key, see "Using the Encryption Key" on page 4-20 and the documentation provided with your TACACS+ server application.

You can enter up to three IP addresses; one first-choice and two (optional) backups (one second-choice and one third-choice).
Use **show tacacs** to view the current IP address list.
If the first-choice TACACS+ server fails to respond to a request, the switch tries the second address, if any, in the show tacacs list. If the second address also fails, then the switch tries the third address, if any.
(See figure 4-3, "Example of the Switch's TACACS+ Configuration Listing" on 4-8.)

The priority (first-choice, second-choice, and third-choice) of a TACACS+ server in the switch's TACACS+ configuration depends on the order in which you enter the server IP addresses:

1. When there are no TACACS+ servers configured, entering a server IP address makes that server the first-choice TACACS+ server.

2. When there is one TACACS+ serves already configured, entering another server IP address makes that server the second-choice (backup) TACACS+ server.

3. When there are two TACACS+ servers already configured, entering another server IP address makes that server the third-choice (backup) TACACS+ server.

- The above position assignments are fixed. Thus, if you remove one server and replace it with another, the new server assumes the priority position that the removed server had. For example, suppose you configured three servers, A, B, and C, configured in order:
First-Choice:A
Second-Choice:B
Third-Choice: C

- If you removed server B and then entered server X, the TACACS+ server order of priority would be:
First-Choice:A
Second-Choice:X
Third-Choice: C

- If there are two or more vacant slots in the TACACS+ server priority list and you enter a new IP address, the new address will take the vacant slot with the highest priority. Thus, if A, B, and C are configured as above and you (1) remove A and B, and (2) enter X and Y (in that order), then the new TACACS+ server priority list would be X, Y, and C.

- The easiest way to change the order of the TACACS+ servers in the priority list is to remove all server addresses in the list and then re-enter them in order, with the new first-choice server address first, and so on.

To add a new address to the list when there are already three addresses present, you must first remove one of the currently listed addresses.

See also "General Authentication Process Using a TACACS+ Server" on page 4-18.

Name	Default	Range
key *<key-string>*	none (null)	n/a

Specifies the optional, global "encryption key" that is also assigned in the TACACS+ server(s) that the switch will access for authentication. This option is subordinate to any "per-server" encryption keys you assign, and applies only to accessing TACACS+ servers for which you have not given the switch a "per-server" key. (See the **host *<ip-addr>* [key *<key-string>*** entry at the beginning of this table.)

For more on the encryption key, see "Using the Encryption Key" on page 4-20 and the documentation provided with your TACACS+ server application.

timeout <1 - 255>	5 sec	1 - 255 sec

Specifies how long the switch waits for a TACACS+ server to respond to an authentication request. If the switch does not detect a response within the timeout period, it initiates a new request to the next TACACS+ server in the list. If all TACACS+ servers in the list fail to respond within the timeout period, the switch uses either local authentication (if configured) or denies access (if **none** configured for local authentication).

Adding, Removing, or Changing the Priority of a TACACS+ Server.

Suppose that the switch was already configured to use TACACS+ servers at 10.28.227.10 and 10.28.227.15. In this case, 10.28.227.15 was entered first, and so is listed as the first-choice server:

```
HP Switch# show tacacs
 Status and Counters - TACACS Information
  Timeout : 5
  Encryption Key : paris-1          First-Choice TACACS+ Server

  Server IP Addr   Opens  Closes Ahorts Errors Pkts Rx Pkts Tx
  --------------   ------ ------ ------ ------ ------- -------
  10.28.227.15     0      0      0      0      0       0
  10.28.227.15     0      0      0      0      0       0
```

Figure 4-6. Example of the Switch with Two TACACS+ Server Addresses Configured

To move the "first-choice" status from the "15" server to the "10" server, use the **no tacacs-server host *<ip-addr>*** command to delete both servers, then use **tacacs-server host *<ip-addr>*** to re-enter the "10" server first, then the "15" server.

The servers would then be listed with the new "first-choice" server, that is:

```
HP Switch# show tacacs
 Status and Counters - TACACS Information
  Timeout : 5
  Encryption Key :        The "10" server is now the "first-choice" TACACS+ authentication device.

  Server IP Addr   Opens  Closes Ahorts Errors Pkts Rx Pkts Tx
  --------------   ------ ------ ------ ------ ------- -------
  10.28.227.10     0      0      0      0      0       0
  10.28.227.15     0      0      0      0      0       0
```

Figure 4-7. Example of the Switch After Assigning a Different "First-Choice" Server

To remove the 10.28.227.15 device as a TACACS+ server, you would use this command:

```
HP Switch(config)# no tacacs-server host 10.28.227.15
```

Configuring an Encryption Key. Use an encryption key in the switch if the switch will be requesting authentication from a TACACS+ server that also uses an encryption key. (If the server expects a key, but the switch either does not provide one, or provides an incorrect key, then the authentication attempt will fail.) Use a *global encryption key* if the same key applies to all TACACS+ servers the switch may use for authentication attempts. Use a *per-server encryption key* if different servers the switch may use will have different keys. (For more details on encryption keys, see "Using the Encryption Key" on page 4-20.)

To configure **north01** as a global encryption key:

```
HP Switch(config) tacacs-server key north01
```

To configure **north01** as a per-server encryption key:

```
HP Switch(config)# tacacs-server host 10.28.227.63 key
north01
```

An encryption key can contain up to 100 characters, without spaces, and is likely to be case-sensitive in most TACACS+ server applications.

To delete a global encryption key from the switch, use this command:

```
HP Switch(config)# no tacacs-server key
```

To delete a per-server encryption key in the switch, re-enter the tacacs-server host command without the key parameter. For example, if you have **north01** configured as the encryption key for a TACACS+ server with an IP address of 10.28.227.104 and you want to eliminate the key, you would use this command:

```
HP Switch(config)# tacacs-server host 10.28.227.104
```

Note
You can save the encryption key in a configuration file by entering this command:

> **HP Switch(config)# tacacs-server key <keystring>**

The *<keystring>* parameter is the encryption key in clear text.

Note   The **show tacacs** command lists the global encryption key, if configured. However, to view any configured per-server encryption keys, you must use **show config** or **show config running** (if you have made TACACS+ configuration changes without executing **write mem**).

Configuring the Timeout Period. The timeout period specifies how long the switch waits for a response to an authentication request from a TACACS+ server before either sending a new request to the next server in the switch's Server IP Address list or using the local authentication option. For example, to change the timeout period from 5 seconds (the default) to 3 seconds:

```
HP Switch(config)# tacacs-server timeout 3
```

How Authentication Operates

General Authentication Process Using a TACACS+ Server

Authentication through a TACACS+ server operates generally as described below. For specific operating details, refer to the documentation you received with your TACACS+ server application.
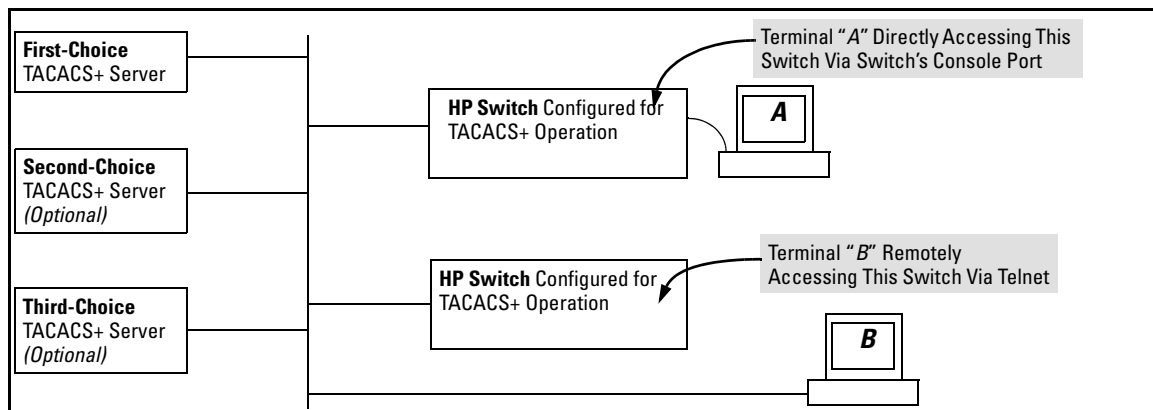
Figure 4-8. Using a TACACS+ Server for Authentication

Using figure 4-8, above, after either switch detects an operator's logon request from a remote or directly connected terminal, the following events occur:

1. The switch queries the first-choice TACACS+ server for authentication of the request.

    • If the switch does not receive a response from the first-choice TACACS+ server, it attempts to query a secondary server. If the switch does not receive a response from any TACACS+ server, then it uses its own local username/password pairs to authenticate the logon request. (See "Local Authentication Process" on page 4-19.)

    • If a TACACS+ server recognizes the switch, it forwards a username prompt to the requesting terminal via the switch.

2. When the requesting terminal responds to the prompt with a username, the switch forwards it to the TACACS+ server.

3. After the server receives the username input, the requesting terminal receives a password prompt from the server via the switch.

4. When the requesting terminal responds to the prompt with a password, the switch forwards it to the TACACS+ server and one of the following actions occurs:

    • If the username/password pair received from the requesting terminal matches a username/password pair previously stored in the server, then the server passes access permission through the switch to the terminal.

    • If the username/password pair entered at the requesting terminal does not match a username/password pair previously stored in the server, access is denied. In this case, the terminal is again prompted to enter a username and repeat steps 2 through 4. In the default configuration, the switch allows up to three attempts to authenticate a login session. If the requesting terminal exhausts the attempt limit without a successful TACACS+ authentication, the login session is terminated and the operator at the requesting terminal must initiate a new session before trying again.

Local Authentication Process

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

■ "Local" is the authentication option for the access method being used.

■ TACACS+ is the primary authentication mode for the access method being used. However, the switch was unable to connect to any TACACS+ servers (or no servers were configured) AND **Local** is the secondary authentication mode being used.

(For a listing of authentication options, see table 4-2, "Primary/Secondary Authentication Table" on 4-12.)

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

■ If the operator at the requesting terminal correctly enters the username/password pair for either access level, access is granted.

■ If the username/password pair entered at the requesting terminal does not match either username/password pair previously configured locally in the switch, access is denied. In this case, the terminal is again prompted to enter a username/password pair. In the default configuration, the switch allows up to three attempts. If the requesting terminal exhausts the attempt limit without a successful authentication, the login session is terminated and the operator at the requesting terminal must initiate a new session before trying again.

Note

The switch's menu allows you to configure only the local Operator and Manager passwords, and not any usernames. In this case, all prompts for local authentication will request only a local password. However, if you use the CLI or the web browser interface to configure usernames for local access, you will see a prompt for both a local username and a local password during local authentication.

Using the Encryption Key

General Operation

When used, the encryption key (sometimes termed "key", "secret key", or "secret") helps to prevent unauthorized intruders on the network from reading username and password information in TACACS+ packets moving between the switch and a TACACS+ server. At the TACACS+ server, a key may include both of the following:

■ **Global key:** A general key assignment in the TACACS+ server application that applies to all TACACS-aware devices for which an individual key has not been configured.

■ **Server-Specific key:** A unique key assignment in the TACACS+ server application that applies to a specific TACACS-aware device.

Note

Configure a key in the switch only if the TACACS+ server application has this exact same key configured for the switch. That is, if the key parameter in switch "X" does not exactly match the key setting for switch "X" in the TACACS+ server application, then communication between the switch and the TACACS+ server will fail.

Thus, on the TACACS+ server side, you have a choice as to how to implement a key. On the switch side, it is necessary only to enter the key parameter so that it exactly matches its counterpart in the server. For information on how to configure a general or individual key in the TACACS+ server, refer to the documentation you received with the application.

Encryption Options in the Switch

When configured, the encryption key causes the switch to encrypt the TACACS+ packets it sends to the server. When left at "null", the TACACS+ packets are sent in clear text. The encryption key (or just "key") you configure in the switch must be identical to the encryption key configured in the corresponding TACACS+ server. If the key is the same for all TACACS+ servers the switch will use for authentication, then configure a global key in the switch. If the key is different for one or more of these servers, use "server-specific" keys in the switch. (If you configure both a global key and one or more per-server keys, the per-server keys will override the global key for the specified servers.)

For example, you would use the next command to configure a global encryption key in the switch to match a key entered as **north40campus** in two target TACACS+ servers. (That is, both servers use the same key for your switch.) Note that you do not need the server IP addresses to configure a global key in the switch:

```
HP Switch(config)# tacacs-server key north40campus
```

Suppose that you subsequently add a third TACACS+ server (with an IP address of 10.28.227.87) that has **south10campus** for an encryption key. Because this key is different than the one used for the two servers in the previous example, you will need to assign a server-specific key in the switch that applies only to the designated server:

```
HP Switch(config)# tacacs-server host 10.28.227.87 key
south10campus
```

With both of the above keys configured in the switch, the **south10campus** key overrides the **north40campus** key only when the switch tries to access the TACACS+ server having the 10.28.227.87 address.

Controlling Web Browser Interface Access When Using TACACS+ Authentication

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

- Configure local authentication (a Manager user name and password and, optionally, an Operator user name and password) on the switch.

- Configure the switch's Authorized IP Manager feature to allow web browser access only from authorized management stations. (The Authorized IP Manager feature does not interfere with TACACS+ operation.)

- Disable web browser access to the switch by going to the System Information screen in the Menu interface and configuring the **Web Agent Enabled** parameter to **No**.

Messages Related to TACACS+ Operation

The switch generates the CLI messages listed below. However, you may see other messages generated in your TACACS+ server application. For information on such messages, refer to the documentation you received with the application.

CLI Message	Meaning
Connecting to Tacacs server	The switch is attempting to contact the TACACS+ server identified in the switch's **tacacs-server** configuration as the **first-choice** (or only) TACACS+ server.
Connecting to secondary Tacacs server	The switch was not able to contact the **first-choice** TACACS+ server, and is now attempting to contact the next (secondary) TACACS+ server identified in the switch's **tacacs-server** configuration.
Invalid password	The system does not recognize the username or the password or both. Depending on the authentication method (**tacacs** or **local**), either the TACACS+ server application did not recognize the username/password pair or the username/password pair did not match the username/password pair configured in the switch.

CLI Message	Meaning
No Tacacs servers responding	The switch has not been able to contact any designated TACACS+ servers. If this message is followed by the **Username** prompt, the switch is attempting local authentication.
Not legal combination of authentication methods	**For console access,** if you select **tacacs** as the primary authentication method, you must select **local** as the secondary authentication method. This prevents you from being locked out of the switch if all designated TACACS+ servers are inaccessible to the switch.
Record already exists	When resulting from a **tacacs-server host** <*ip addr*> command, indicates an attempt to enter a duplicate TACACS+ server IP address.

Operating Notes

■ If you configure Authorized IP Managers on the switch, it is not necessary to include any devices used as TACACS+ servers in the authorized manager list. That is, authentication traffic between a TACACS+ server and the switch is not subject to Authorized IP Manager controls configured on the switch. Also, the switch does not attempt TACACS+ authentication for a management station that the Authorized IP Manager list excludes because, independent of TACACS+, the switch already denies access to such stations.

■ When TACACS+ is not enabled on the switch—or when the switch's only designated TACACS+ servers are not accessible— setting a local Operator password without also setting a local Manager password does not protect the switch from manager-level access by unauthorized persons.

■ When using the **copy** command to transfer a configuration to a TFTP server, any optional, server-specific and global encryption keys (page 4-13) in the TACACS configuration will not be included in the transferred file. Otherwise, a security breach could occur, allowing access to the TACACS+ username/password information.

RADIUS Authentication, Authorization and Accounting

Overview

Feature	Default	Menu	CLI	Web
Configuring RADIUS Authentication	None	n/a	5-5	n/a
Configuring RADIUS Accounting	None	n/a	5-29	n/a
Viewing RADIUS Statistics	n/a	n/a	5-36	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 HP switch:

■   Serial port (Console)

■   Telnet

■   SSH

■   Web

■   Port-Access

Note   The switch does not support RADIUS security for SNMP (network management) access. For information on blocking unauthorized access through the Web browser interface, refer to "Controlling Web Browser Interface Access When Using RADIUS Authentication" on page 5-20.

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.

SNMP Access to the Switch's Authentication Configuration MIB

The switch's default configuration allows SNMP access to the hpSwitchAuth MIB (Management Information Base). A management station running an SNMP networked device management application such as IMC or HP Open-View can access the switch's MIB for read access to the switch's status and read/write access to the switch's configuration. For more information, including the CLI command to use for disabling this feature, refer to "Using SNMP To View and Configure Switch Authentication Features" on page 5-17.

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, the HP switch 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 HP switch, 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 **show radius** (page 5-36). 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 5-40.)

■ 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 HP switch, 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 configuring the switch, collect the information outlined below.

Table 5-1.   Preparation for Configuring RADIUS on the Switch

- Determine the access methods (console, Telnet, Port-Access (802.1X), SSH, and/or Web browser interface) 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.

```
Switch> show authentication

 Status and Counters - Authentication Information

  Login Attempts : 3
  Respect Privilege : Disabled

               | Login      Login      Enable     Enable
  Access Task  | Primary    Secondary  Primary    Secondary
  ----------- + ---------- ---------- ---------- ---------
  Console      | Radius     Local      Radius     Local
  Telnet       | Radius     None       Radius     None
  Port-Access  | EapRadius
  Webui        | Radius     None       Radius     None
  SSH          | Radius     None       Radius     None
  Web-Auth     | ChapRadius
  MAC-Auth     | ChapRadius
```

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 5-1. 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	Page
aaa authentication	5-7
< console | telnet | ssh | web > < enable | login <local | radius>> web-based | mac-based <chap-radius | peap-radius>	5-7
< local | none | authorized >	5-7
[login privilege-mode]	5-9
[no] radius-server host < *IP-address* >	5-11
[auth-port < *port-number* >]	5-11
[acct-port < *port-number* >]	5-11, 5-32
[key < *server-specific key-string* >]	5-11
[no] radius-server key < *global key-string* >	5-11
radius-server timeout < 1 - 15>	5-14
radius-server retransmit < 1 - 5 >	5-14
[no] radius-server dead-time < 1 - 1440 >	5-15
show radius	5-36
[< host < *ip-address*>]	5-36
show authentication	5-38
show radius authentication	5-38

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
   - Web browser interface
   - 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:** Number of retries when there is no server response to RADIUS authentication request. (Default: 3; range 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 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 5-29.)

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 6-1.

■ **Web:** Web browser interface.

You can also use RADIUS for Port-Based Access authentication. Refer to "Configuring Port-Based and Client-Based Access Control (802.1X)" on page 10-1.

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 | web | < enable | login <local | radius>> web-based | mac-based <chap-radius | peap-radius>>

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

<console | telnet | ssh | web>

[< local | none | authorized >]

Provides 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 locked out of the switch in the event of a failure in other access methods.*

<<web-based|mac-based>login> <chap-radius|peap-mschapv2>:
*Password authentication for web-based or mac-based port
access to the switch. Use* **peap-mschapv2** *when you want pass-
word verification without requiring access to a plain text
password; it is more secure.*
Default: **chap-radius**

[none | authorized]: *Provides options for secondary
authentication. The* **none** *option specifies that a backup
authentication method is not used. The* **authorized**
option allows access without authentication. (default:
none*).*

In certain situations, RADIUS servers can become isolated from the network.
Users are not able to access the network resources configured with RADIUS
access protection and are rejected. To address this situation, configuring the
authorized secondary authentication method allows users unconditional
access to the network when the primary authentication method fails because
the RADIUS servers are unreachable.

Caution

Configuring **authorized** as the secondary authentication method used when
there is a failure accessing the RADIUS servers allows clients to access the
network unconditionally. Use this method with care.

Figure 5-2 shows an example of the **show authentication** command displaying
authorized as the secondary authentication method for port-access, Web-auth
access, and MAC-auth access. Since the configuration of **authorized** means no
authentication will be performed and the client has unconditional access to
the network, the "Enable Primary" and "Enable Secondary" fields are not
applicable (N/A).

```
HP Switch(config)# show authentication
 Status and Counters - Authentication Information
  Login Attempts : 3
  Respect Privilege : Disabled
              | Login      Login      Enable     Enable
  Access Task | Primary    Secondary  Primary    Secondary
  ----------- + ---------- ---------- ---------- ----------
  Console     | Local      None       Local      None
  Telnet      | Local      None       Local      None
  Port-Access | Local      Authorized N/A        N/A
  Webui       | Local      None       Local      None
  SSH         | Local      None       Local      None
  Web-Auth    | ChapRadius Authorized N/A        N/A
  MAC-Auth    | ChapRadius Authorized N/A        N/A
```

The access methods
with secondary
authentication
configured as **authorized**
allows the client access
to the network even if the
RADIUS server is
unreachable.

Figure 5-2. Example of AAA Authentication Using Authorized for the Secondary Authentication Method

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):

```
HP Switch(config)# aaa authentication telnet login radius none
HP Switch(config)# aaa authentication telnet enable radius none
HP Switch(config)# aaa authentication ssh login radius none
HP Switch(config)# aaa authentication ssh enable radius none
HP Switch(config)# show authentication

 Status and Counters - Authentication Information

  Login Attempts : 3
  Respect Privilege : Disabled


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

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

Figure 5-3. 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 5-19.

2. Enable the (Optional) Access Privilege Option

In the default RADIUS operation, the switch automatically admits any authenticated client to the Login (Operator) privilege level, even if the RADIUS server specifies Enable (Manager) access for that client. Thus, an authenticated user authorized for the Manager privilege level must authenticate again to change privilege levels. Using the optional **login privilege-mode** command overrides this default behavior for clients with Enable (manager) access. That is, with **privilege-mode** enabled, the switch immediately allows Enable (Manager) access to a client for whom the RADIUS server specifies this access level

.

Syntax: [no] aaa authentication login privilege-mode

When enabled, the switch reads the Service-Type field in the client authentication received from a RADIUS server. The following table describes the applicable Service-Type values and corresponding client access levels the switch allows upon authentication by the server.

Service-Type	Value	Client Access Level
Administrative-User	6	Manager
NAS-Prompt-User	7	Operator
Any Other Type	Any Value *Except* 6 or 7	Access Denied

This feature applies to console (serial port), Telnet, SSH, and web browser interface access to the switch. It does not apply to 802.1X port-access.

Notes: *While this option is enabled, a Service-Type value other than 6 or 7, or an unconfigured (null) Service-Type causes the switch to deny access to the requesting client.*

*The **no** form of the command returns the switch to the default RADIUS authentication operation. The default behavior for most interfaces is that a client authorized by the RADIUS server for Enable (Manager) access will be prompted twice, once for Login (Operator) access and once for Enable access. In the default RADIUS authentication operation, the switch's web browser interface requires only one successful authentication request. For more information on configuring the Service Type in your RADIUS application, refer to the documentation provided with the application.*

3. 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.

Note	If you want to configure RADIUS accounting on the switch, go to page 5-29: "Configuring RADIUS Accounting" 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 5-40.)*

[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)*

[dyn-authorization]

>*Enables or disables dynamic authorization control from this host.*

[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 5-4 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".

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

Figure 5-4. Sample Configuration for RADIUS Server Before Changing the Key and Adding Another Server

To make the changes listed prior to figure 5-4, you would do the following:

```
HP   Switch(config)#   radius-server   host   10.33.18.127   key
source0127

HP   Switch(config)#   radius-server   host   10.33.18.119   key
source0119

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

                  Auth   Acct
  Server IP Addr  Port  Port  Encryption Key
  --------------- ----- ----- ------------------
  10.33.18.127    1812  1813  source0127
  10.33.18.119    1812  1813  source0119
```

Changes the key for the existing server to "source0127"

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

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

Figure 5-5. 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 5-40.

4. 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 "3. Configure the Switch To Access a RADIUS Server" on page 5-11.)

- **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.

- **Change of Authorization port:** The **dyn-autz-port** parameter specifies the UDP port number that listens for the Change of Authorization and Disconnect messages.

Syntax: aaa authentication num-attempts < 1 - 10 >

> *Specifies how many tries for entering the correct user-name 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 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)*

dyn-autz-port <1024-49151>

> *Specifies the UDP port number that listens for a Change of Authorization or Disconnect messages.*

> *Default: 3799*

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, refer to "RADIUS-Related Problems" in the Troubleshooting chapter of the *Management and Configuration Guide* for your switch.

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.

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

Figure 5-6. Example of Global Configuration Exercise for RADIUS Authentication

```
HP Switch# show authentication

 Status and Counters - Authentication Information

  Login Attempts : 2
  Respect Privilege : Disabled

               | Login        Login        Enable       Enable
  Access Task  | Primary      Secondary    Primary      Secondary
  ----------- + ----------   ----------   ----------   ----------
  Console      | Local        None         Local        None
  Telnet       | Radius       None         Radius       None
  Port-Access  | Local
  Webui        | Local        None         Local        None
  SSH          | Radius       None         Radius       None
  Web-Auth     | ChapRadius
  MAC-Auth     | ChapRadius


HP Switch# show radius

 Status and Counters - General RADIUS Information

  Deadtime(min) : 5
  Timeout(secs) : 3
  Retransmit Attempts : 2
  Global Encryption Key : My-Global-Key-1099

                 Auth  Acct
  Server IP Addr Port  Port  Encryption Key
  -------------- ----- ----- ---------------------------
  10.33.18.127   1812  1813  source0127
  10.33.18.119   1812  1813
  10.33.18.151   1812  1813
```

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

Global RADIUS parameters from figure 5-6.

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

These two servers will use the global encryption key.

Figure 5-7. Listings of Global RADIUS Parameters Configured In Figure 5-6

Using SNMP To View and Configure Switch Authentication Features

SNMP MIB object access is available for switch authentication configuration (hpSwitchAuth) features. This means that the switches covered by this Guide allow, by default, manager-only SNMP read/write access to a subset of the authentication MIB objects for the following features:

■　number of primary and secondary login and enable attempts

■　TACACS+ server configuration and status

■　RADIUS server configuration

■　 selected 802.1X settings

■　key management subsystem chain configuration

■　key management subsystem key configuration

■　OSPF interface authentication configuration

■　local switch operator and manager usernames and passwords

With SNMP access to the hpSwitchAuth MIB enabled, a device with management access to the switch can view the configuration for the authentication features listed above (excluding usernames, passwords, and keys). Using SNMP sets, a management device can change the authentication configuration (*including* changes to usernames, passwords, and keys). Operator read/write access to the authentication MIB is always denied.

Security Notes　All usernames, passwords, and keys configured in the hpSwitchAuth MIB are not returned via SNMP, and the response to SNMP queries for such information is a null string. However, SNMP sets can be used to configure username, password, and key MIB objects.

To help prevent unauthorized access to the switch's authentication MIB, HP recommends following the "SNMP Security Guidelines" on page 1-10.

If you do not want to use SNMP access to the switch's authentication configuration MIB, then use the **snmp-server mib hpswitchauthmib excluded** command to disable this access, as described in the next section.

If you choose to leave SNMP access to the security MIB open (the default setting), HP recommends that you configure the switch with the SNMP version 3 management and access security feature, and disable SNMP version 2c access. (Refer to "Access Security Features" on page 1-2.)

Changing and Viewing the SNMP Access Configuration

Syntax: snmp-server mib hpswitchauthmib < excluded | included >

> **included:** *Enables manager-level SNMP read/write access to the switch's authentication configuration (hpSwitchAuth) MIB.*
> **excluded:** *Disables manager-level SNMP read/write access to the switch's authentication configuration (hpSwitchAuth) MIB.*
> *(Default:* **included** *)*

Syntax: show snmp-server

> *The output for this command has been enhanced to display the current access status of the switch's authentication configuration MIB in the* **Excluded MIBs** *field.*

For example, to disable SNMP access to the switch's authentication MIB and then display the result in the Excluded MIB field, you would execute the following two commands.

```
HP Switch(config)# snmp-server mib hpswitchauthmib excluded
HP Switch(config)# show snmp-server

 SNMP Communities

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

 Trap Receivers

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

  Send Authentication Traps [No] : No

  Address                 Community         Events Sent in Trap
  --------------------- --------------- -------------------

 Excluded MIBs

  hpSwitchAuthenticationMIB
```

This command disables SNMP security MIB access.

Indicates that SNMP security MIB access is disabled, which is the nondefault setting.

Figure 5-8. Disabling SNMP Access to the Authentication MIB and Displaying the Result

An alternate method of determining the current Authentication MIB access state is to use the **show run** command.

```
HP Switch(config)# show run

Running configuration:

; J9229A Configuration Editor; Created on release #YA.15.XX

hostname "Switch"
snmp-server mib hpSwitchAuthMIB excluded
ip default-gateway 10.10.24.55
snmp-server community "public" Operator
vlan 1
    name "DEFAULT_VLAN"
    untagged 1-24
    ip address 10.10.24.100 255.255.255.0
    exit
password manager
```

Indicates that SNMP access to the authentication configuration MIB (hpSwitchAuth) is disabled.
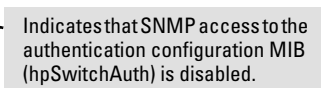
Figure 5-9. Using the show run Command to View the Current Authentication MIB Access State

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

To help prevent unauthorized access through the web browser interface, do one or more of the following:

■ Configure the switch to support RADIUS authentication for web browser interface access (See chapter 3, "Web and MAC Authentication").

■ Options for the switches covered in this guide:

• 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.)

• Use one of the following methods to disable web browser access to the switch via http (Port 80):

CLI: **no web-management**

Menu Interface—From the Main menu, select the following:

2. Switch Configuration

1. System Information

Web Agent Enabled: No

Commands Authorization

The RADIUS protocol combines user authentication and authorization steps into one phase. The user must be successfully authenticated before the RADIUS server will send authorization information (from the user's profile) to the Network Access Server (NAS). After user authentication has occurred, the authorization information provided by the RADIUS server is stored on the NAS for the duration of the user's session. Changes in the user's authorization profile during this time will not be effective until after the next authentication occurs.

You can limit the services for a user by enabling AAA RADIUS authorization. The NAS uses the information set up on the RADIUS server to control the user's access to CLI commands.

The authorization type implemented on the switches covered in this guide is the "commands" method. This method explicitly specifies on the RADIUS server which commands are allowed on the client device for authenticated users. This is done on a per-user or per-group basis.

N o t e   The commands authorization will only be executed for commands entered from Telnet, SSH, or console sessions. The Web management interface is not supported.

By default, all users may execute a minimal set of commands regardless of their authorization status, for example, "exit" and "logout". This minimal set of commands can prevent deadlock on the switch due to an error in the user's authorization profile on the RADIUS server.

Enabling Authorization

To configure authorization for controlling access to the CLI commands, enter this command at the CLI.

Syntax: [no] aaa authorization <commands> <radius | none>

> *Configures authorization for controlling access to CLI commands. When enabled, the switch checks the list of commands supplied by the RADIUS server during user authentication to determine if a command entered by the user can be executed.*

radius: *The NAS requests authorization information from the RADIUS server. Authorization rights are assigned by user or group.*

none: *The NAS does not request authorization information.*

For example, to enable the RADIUS protocol as the authorization method:

```
HP Switch(config)# aaa authorization commands radius
```

When the NAS sends the RADIUS server a valid username and password, the RADIUS server sends an Access-Accept packet that contains two attributes—the command list and the command exception flag. When an authenticated user enters a command on the switch, the switch examines the list of commands delivered in the RADIUS Access-Accept packet as well as the command exception flag, which indicates whether the user has permission to execute the commands in the list. See *Configuring the RADIUS Server* on page 5-23.

After the Access-Accept packet is deliver, the command list resides on the switch. Any changes to the user's command list on the RADIUS server are not seen until the user is authenticated again.

Displaying Authorization Information

You can show the authorization information by entering this command:

Syntax: show authorization

> *Configures authorization for controlling access to CLI commands. When enabled, the switch checks the list of commands supplied by the RADIUS server during user authentication to determine if a command entered by the user can be executed.*

An example of the output is shown.

```
HP Switch(config)# show authorization

 Status and Counters - Authorization Information

  Type     | Method

  -------- + ------

  Commands | RADIUS
```

Figure 5-10. Example of Show Authorization Command

Configuring Commands Authorization on a RADIUS Server

Using Vendor Specific Attributes (VSAs)

Some RADIUS-based features implemented on HP switches use HP VSAs for information exchange with the RADIUS server. RADIUS Access-Accept packets sent to the switch may contain the vendor-specific information. The attributes supported with **commands** authorization are:

■   HP-Command-String: List of commands (regular expressions) that are permitted (or denied) execution by the user. The commands are delimited by semi-colons and must be between 1 and 249 characters in length. Multiple instances of this attribute may be present in Access-Accept packets. (A single instance may be present in Accounting-Request packets.)

■   HP-Command-Exception: A flag that specifies whether the commands indicated by the HP-Command-String attribute are permitted or denied to the user. A zero (0) means permit all listed commands and deny all others; a one (1) means deny all listed commands and permit all others.

The results of using the HP-Command-String and HP-Command-Exception attributes in various combinations are shown below.

HP-Command-String	HP-Command-Exception	Description
Not present	Not present	If command authorization is enabled and the RADIUS server does not provide any authorization attributes in an Access-Accept packet, the user is denied access to the server. This message appears: "Access denied: no user's authorization info supplied by the RADIUS server."
Not present	DenyList-PermitOthers(1)	Authenticated user is allowed to execute all commands available on the switch.
Not present	PermitList-DenyOthers(0)	Authenticated user can only execute a minimal set of commands (those that are available by default to any user).
Commands List	DenyList-PermitOthers(1)	Authenticated user may execute all commands except those in the Commands list.
Commands List	PermitList-DenyOthers(0)	Authenticated user can execute only those commands provided in the Commands List, plus the default commands.

HP-Command-String	HP-Command-Exception	Description
Commands List	Not present	Authenticated user can only execute commands from the Commands List, plus the default commands.
Empty Commands List	Not present	Authenticate user can only execute a minimal set of commands (those that are available by default to any user).
Empty Commands List	DenyList-PermitOthers(1)	Authenticated user is allowed to execute all commands available on the switch.
Empty Commands List	PermitList-DenyOthers(0)	Authenticate user can only execute a minimal set of commands (those that are available by default to any user).

You must configure the RADIUS server to provide support for the HP VSAs. There are multiple RADIUS server applications; the two examples below show how a dictionary file can be created to define the VSAs for that RADIUS server application.

Example Configuration on Cisco Secure ACS for MS Windows

It is necessary to create a dictionary file that defines the VSAs so that the RADIUS server application can determine which VSAs to add to its user interface. The VSAs will appear below the standard attributes that can be configured in the application.

The dictionary file must be placed in the proper directory on the RADIUS server. Follow these steps.

1. Create a dictionary file (for example, hp.ini) containing the HP VSA definitions, as shown in the example below.

```
;[User Defined Vendor]
;
; The Name and IETF vendor code and any VSAs MUST be unique.
;
; One or more VSAs named (max 255)
;
; Each named VSA requires a definition section…
;
; Types are STRING, INTEGER, IPADDR
;
; The profile specifies usage, IN for accounting, OUT for
  authorization,
```

```
; MULTI if more than a single instance is allowed per
  RADIUS message.
; Combinations are allowed, e.g. "IN", "MULTI OUT",
  "MULT IN OUT"
;
; Enumerations are optional for INTEGER attribute types

[User Defined Vendor]

Name=HP
IETF Code=11
VSA 2=Hp-Command-String
VSA 3=Hp-Command-Exception

[Hp-Command-String]

Type=STRING
Profile=IN OUT

[Hp-Command-Exception]

Type=INTEGER
Profile=IN OUT

Enums=Hp-Command-Exception-Types

[Hp-Command-Exception-Types]

0=PermitList
1=DenyList
```

2. Copy the hp.ini dictionary file to c:\program files\cisco acs 3.2\utils (or
    the \utils directory wherever acs is installed).

3. From the command prompt execute the following command:

    ```
    c:\Program files\CiscoSecure ACS v3.2\utils>
    csutil -addudv 0 hp.ini
    ```

    The zero (0) is the slot number. You will see some processing messages:

    ```
    Adding or removing vendors requires ACS services to be
    re-started. Please make sure regedit is not running as
    it can prevent registry backup/restore operations.

    Are you sure you want to proceed?  (Y or N) y
    ```

```
Parsing [.\hp.ini] for addition at UDV slot [0]
Stopping any running services
Creating backup of current config

Adding Vendor [HP} added as [RADIUS (HP)]

Done
Checking new configuration...
New configuration OK
Re-starting stopped services
```

4. Start the registry editor (regedit) and browse to HKEY_LOCAL_MACHINE\software\cisco\CiscoAAA v3.2\NAS Vendors tree.

   Cisco adds the entry into this tree for each custom vendor. The id is 100 + the slot number used in the previous command (100 + 0, as it was added in slot 0). Look in the key to verify the vendor name and id.

5. Go to:

   HKEY_LOCAL_MACHINE\software\cisco\CiscoAAAv3.2\
   CSRadius\ExtensionPoints\002\AssociatedWithVendors

6. Right click and then select **New > key**. Add the vendor Id number that you determined in step 4 (100 in the example).

7. Restart all Cisco services.

8. The newly created HP RADIUS VSA appears only when you configure an AAA client (NAS) to use the HP VSA RADIUS attributes. Select Network Configuration and add (or modify) an AAA entry. In the Authenticate Using field choose RADIUS(HP) as an option for the type of security control protocol.

9. Select **Submit + Restar**t to effect the change. The HP RADIUS VSA attributes will appear in Cisco ACS configurations, for example, "Interface Configuration", "Group Setup", "User Setup".

To enable the processing of the HP-Command-String VSA for RADIUS accounting:

1. Select **System Configuration**.

2. Select **Logging**.

3. Select **CSV RADIUS Accounting**. In the Select Columns to Log section, add the HP-Command-String attribute to the Logged Attributes list.

4. Select **Submit**.

5. Select **Network Configuration**. In the AAA Clients section, select an entry in the AAA Client Hostname column. You will go to the AAA Client Setup screen.

6. Check the box for **Log Update/Watchdog Packets from this AAA Client**.

7. Click **Submit + Restart**. You should be able to see the HP-Command-String attribute in the RADIUS accounting reports.

You can enter the commands you wish to allow or deny with the special characters used in standard regular expressions (c, ., \, [list], [^list], *, ^, $). Commands must be between 1-249 characters in length.

Example Configuration Using FreeRADIUS

1. Create a dictionary file (for example, dictionary.hp) containing HP VSA definitions. An example file is:

```
#   dictionary.hp
#
#   As posted to the list by User <user_email>
#
#   Version: $Id: dictionary.hp, v 1.0 2006/02/23 17:07:07
#


VENDOR          Hp      11

# HP Extensions


ATTRIBUTE       Hp-Command-String      2      string   Hp
ATTRIBUTE       Hp-Command-Exception   3      integer  Hp

# Hp-Command-Exception Attribute Values

VALUE           Hp-Command-Exception    Permit-List      0
VALUE           Hp-Command-Exception    Deny-List        1
```

2. Find the location of the dictionary files used by FreeRADIUS (try /usr/local/share/freeradius).

3. Copy dictionary.hp to that location. Open the existing dictionary file and add this entry:

   $ INCLUDE dictionary.hp

4. You can now use HP VSAs with other attributes when configuring user entries.

Additional RADIUS Attributes

The following attributes are included in Access-Request and Access-Accounting packets sent from the switch to the RADIUS server to advertise switch capabilities, report information on authentication sessions, and dynamically reconfigure authentication parameters:

■ MS-RAS-Vendor (RFC 2548): Allows HP switches to inform a Microsoft RADIUS server that the switches are from HP. This feature assists the RADIUS server in its network configuration.

■ HP-capability-advert: HP proprietary RADIUS attribute that allows a switch to advertise its current capabilities to the RADIUS server for port-based (MAC, Web, or 802.1X) authentication; for example, HP VSAs for port QoS, ingress rate-limiting, IDM filter rules, RFC 4675 QoS and VLAN attributes, and RFC 3580 VLAN-related attributes.

The RADIUS server uses this information to make a more intelligent policy decision on the configuration settings to return to the switch for a client session.

■ HP-acct-terminate-cause: HP proprietary RADIUS accounting attribute that allows a switch to report to the RADIUS server why an authentication session was terminated. This information allows customers to diagnose network operational problems and generate reports on terminated sessions. This attribute provides extended information on the statistics provided by the acct-terminate-cause attribute.

■ change-of-authorization (RFC 3576: Dynamic Authorization Extensions to RADIUS): A mechanism that allows a RADIUS server to dynamically terminate or change the authorization parameters (such as VLAN assignment) used in an active client session on the switch. The switch (NAS) does not have to initiate the exchange.

For example, for security reasons you may want to limit the network services granted to an authenticated user. In this case, you can change the user profile on the RADIUS server and have the new authorization settings take effect immediately in the active client session. The change-of-authorization attribute provides the mechanism to dynamically update an active client session with a new user policy that is sent in RADIUS packets.

Configuring RADIUS Accounting

RADIUS Accounting Commands	Page
[no] radius-server host < *ip-address* >	5-32
[acct-port < *port-number* >]	5-32
[key < *key-string* >]	5-32
[no] aaa accounting < exec | network | system | commands > start-stop | stop-only> radius	5-34
[no] aaa accounting update periodic < 1 - 525600 > (*in minutes*)	5-35
[no] aaa accounting suppress null-username	5-35
show accounting	5-39
show accounting sessions	5-39
show radius accounting	5-39

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 5-3 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 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-Authentic
- Acct-Delay-Time
- Acct-Input-Octets
- Acct-Input-Packets
- Acct-Output-Octets
- Acct-Output-Packets
- Acct-Session-Id
- Acct-Session-Time
- Acct-Status-Type
- Acct-Terminate-Cause
- Called-Station-Id
- HP-acct-terminate-cause
- MS-RAS-Vendor
- NAS-Identifier
- NAS-IP-Address
- NAS-Port
- Service-Type
- Username

(For 802.1X information for the switch, refer to "Configuring Port-Based and Client-Based Access Control (802.1X)" on page 10-1.)

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

• Acct-Authentic	• Acct-Status-Type	• NAS-Identifier
• Acct-Delay-Time	• Acct-Terminate-Cause	• NAS-IP-Address
• Acct-Session-Id	• Calling-Station-Id	• Service-Type
• Acct-Session-Time	• MS-RAS-Vendor	• Username

■ **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-Authentic	• Acct-Terminate-Cause	• NAS-IP-Address
• Acct-Delay-Time	• Calling-Station-Id	• Service-Type
• Acct-Session-Id	• MS-RAS-Vendor	• Username
• Acct-Session-Time	• NAS-Identifier	

■ **Commands accounting:** Provides records containing information after the execution of a command.

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 5-40.)

■ 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.

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 "3. Configure the Switch To Access a RADIUS Server" on page 5-11.
   - 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 5-11. (Default: null)

2. Configure accounting types and the controls for sending reports to the RADIUS server.

   - **Accounting types:** exec (page 5-30), network (page 5-29), or system (page 5-30)
   - **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 access to 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 5-11. 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 5-11.)

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.).

```
HP Switch(config)# radius-server host 10.33.18.151 acct-port key source0151

HP Switch(config)# write mem

HP Switch(config)# show radius
 Status and Counters - General RADIUS Information
  Deadtime(min) : 5
  Timeout(secs) : 3
  Retransmit Attempts : 2
  Global Encryption Key :

                 Auth  Acct
  Server IP Addr Port  Port  Encryption Key
  -------------- ----- ----- -------------------
  10.33.18.127   1812  1750  source0151
```

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

Figure 5-11. Example of Configuring for a RADIUS Server with a Non-Default Accounting UDP Port Number

The radius-server command as shown in figure 5-11, 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 Accounting Types and the Controls for Sending Reports 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 5-2.)

- **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 time span 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 2.) For information on this feature, refer to "Configuring Port-Based and Client-Based Access Control (802.1X)" on page 10-1.

- **Commands:** When commands authorization is enabled, a record accounting notice is sent after the execution of a command.

Determine how to send accounting data to a RADIUS server:

■ **Start-Stop:**

• Send a start record accounting notice at the beginning of the account-
ing session and a stop record notice at the end of the session. Both
notices include the latest data the switch has collected for the
requested accounting type (Network, Exec, or System).

• Do not wait for an acknowledgement.

The system option (page 5-33) ignores **start-stop** because the switch sends
the accumulated data only when there is a reboot, reload, or accounting
on/off event.

■ **Stop-Only:**

• Send a stop record accounting notice at the end of the accounting
session. The notice includes the latest data the switch has collected
for the requested accounting type (Network, Exec, or System).

• Do not wait for an acknowledgment.

The system option (page 5-33) always delivers **stop-only** operation because
the switch sends the accumulated data only when there is a reboot, reload,
or accounting on/off event.

Syntax: [no] aaa accounting < exec | network | system | commands> < start-stop
| stop-only > radius

*Configures RADIUS accounting type and how data will be
sent to the RADIUS server.*

For example, to configure RADIUS accounting on the switch with **start-stop**
for exec functions and **stop-only** for system functions:

```
HP Switch(config)# aaa accounting exec start-stop radius
HP Switch(config)# aaa accounting system stop-only radius

HP Switch(config)# show accounting
 Status and Counters - Accounting Information
  Interval(min) : 0
  Suppress Empty User : No


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

Configures exec and
system accounting and
controls.

Summarizes the switch's
accounting configuration.

Exec and System accounting are
active. (Assumes the switch is
configured to access a reachable
RADIUS server.)

Figure 5-12. Example of Configuring Accounting Types

3. (Optional) Configure Session Blocking and Interim Updating Options

These optional parameters give you additional control over accounting data.

■ **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)*

Syntax: [no] aaa accounting suppress null-username

Disables accounting for unknown users having no username. (Default: suppression disabled)

To continue the example in figure 5-12, suppose that you wanted the switch to:

■ Send updates every 10 minutes on in-progress accounting sessions.

■ Block accounting for unknown users (no username).

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

HP Switch(config)# show accounting
 Status and Counters - Accounting Information          • Update Period
  Interval(min) : 10
  Suppress Empty User : No                             • Suppress Unknown User

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

Figure 5-13. Example of Optional Accounting Update Period and Accounting Suppression on Unknown User

Viewing RADIUS Statistics

General RADIUS Statistics

Syntax: show radius [host < *ip-addr* >]

> *Shows general RADIUS configuration, including the server IP addresses. Optional form shows data for a specific RADIUS host. To use* **show radius,** *the server's IP address must be configured in the switch, which. requires prior use of the* **radius-server host** *command. (See "Configuring RADIUS Accounting" on page 5-29.)*

```
HP Switch(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 5-14. Example of General RADIUS Information from Show Radius Command

```
HP Switch(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 5-15. RADIUS Server Information From the Show Radius Host Command

Table 5-2.   Values for Show Radius Host Output (Figure 5-15)

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.
Pending Requests	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.
Access Requests	The number of RADIUS Access-Requests the switch has sent since it was last rebooted. (Does not include retransmissions.)
Accounting Requests	The number of RADIUS Accounting-Request packets sent. This does not include retransmissions.
Access Challenges	The number of RADIUS Access-Challenge packets (valid or invalid) received from this server.
Access Accepts	The number of RADIUS Access-Accept packets (valid or invalid) received from this server.
Access Rejects	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 Statistics

Syntax: show authentication

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

show radius authentication

> *Displays NAS identifier and data on the configured RADIUS server and the switch's interactions with this server. (Requires prior use of the **radius-server host** command to configure a RADIUS server IP address in the switch. See "Configuring RADIUS Accounting" on page 5-29.)*

```
HP Switch(config)# show accounting

 Status and Counters - Accounting Information

  Interval(min) : 0
  Suppress Empty User : No

  Type     | Method Mode
  -------- + ------ ----------
  Network  | None
  Exec     | None
  System   | None
  Commands | None
```

Figure 5-16. Example of Login Attempt and Primary/Secondary Authentication Information from the Show Authentication Command

```
HP Switch(config)# show radius authentication
 Status and Counters - RADIUS Authentication Information

  NAS Identifier : HPswitch
  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 5-17. Example of RADIUS Authentication Information from a Specific Server

RADIUS Accounting Statistics

Syntax: show accounting

> *Lists configured accounting interval, "Empty User" suppression status, accounting types, methods, and modes.*

show radius accounting

> *Lists accounting statistics for the RADIUS server(s) configured in the switch (using the* **radius-server host** *command).*

show accounting sessions

> *Lists the accounting sessions currently active on the switch.*

```
HP Switch(config)# show accounting
 Status and Counters - Accounting Information


  Interval(min) : 5
  Suppress Empty User : No


  Type     | Method Mode
  ------- + ------ --------
  Network | None
  Exec    | Radius Start-Stop
  System  | Radius Stop-Only
```
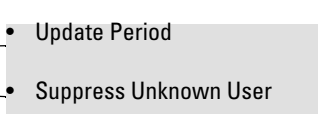
Figure 5-18. Listing the Accounting Configuration in the Switch

```
HP Switch# show radius accounting
 Status and Counters - RADIUS Accounting Information
  NAS Identifier : HPswitch
  Invalid Server Addresses : 0


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

Figure 5-19. Example of RADIUS Accounting Information for a Specific Server

```
  HP Switch# show accounting sessions

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

Figure 5-20. 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:

```
HP Switch# show radius
 Status and Counters — General RADIUS Information
  Deadtime(min)  : 0
  Timeout(secs)  : 5
  Retransmit Attempts : 3
  Global Encryption Key :
                  Auth  Auth
  Server IP Addr  Port  Port  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 10.10.10.1 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 5-21. 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.

```
HP Switch(config)# no radius host 10.10.10.003        Removes the "003" and "001" addresses from
HP Switch(config)# no radius host 10.10.10.001        the RADIUS server list.
HP Switch(config)# radius host 10.10.10.003           Inserts the "003" address in the first position in
HP Switch(config)# radius host 10.10.10.001           the RADIUS server list, and inserts the "001"
                                                       address in the last position in the list.
HP Switch(config)# show radius

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

                 Auth  Auth
  Server IP Addr Port  Port  Encryption Key
  -------------- ----- ----- ----------------
  10.10.10.3      1812  1813                         Shows the new order in which the switch
  10.10.10.2      1812  1813                         searches for a RADIUS server.
  10.10.10.1      1812  1813
```

Figure 5-22. Example of New RADIUS Server Search Order

Configuring Secure Shell (SSH)

Overview

Feature	Default	Menu	CLI	Web
Generating a public/private key pair on the switch	No	n/a	page 6-7	n/a
Using the switch's public key	n/a	n/a	page 6-10	n/a
Enabling SSH	Disabled	n/a	page 6-12	n/a
Enabling client public-key authentication	Disabled	n/a	pages 6-16, 6-19	n/a
Enabling user authentication	Disabled	n/a	page 6-16	n/a

The switches covered in this guide use Secure Shell version 2 (SSHv2) to provide remote access to management functions on the switches via encrypted paths between the switch and management station clients capable of SSH operation.

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.)
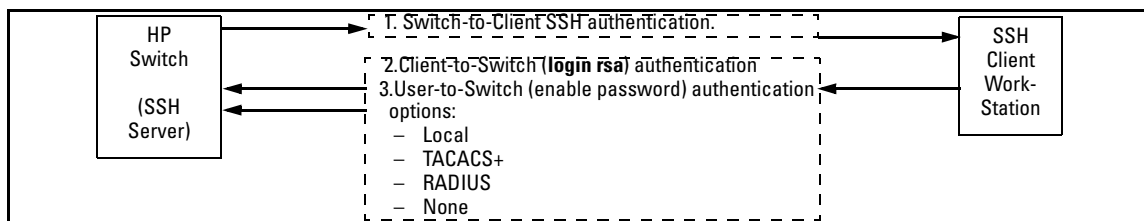
Figure 6-1.   Client Public Key Authentication Model

N o t e   SSH in HP switches is based on the OpenSSH software toolkit. For more information on OpenSSH, visit **www.openssh.com**.

Switch SSH and User Password Authentication . This option is a subset of the client public-key authentication shown in figure 6-1. It occurs if the switch has SSH enabled but does not have login access (**login public-key**) configured to authenticate the client's key. As in figure 6-1, 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.
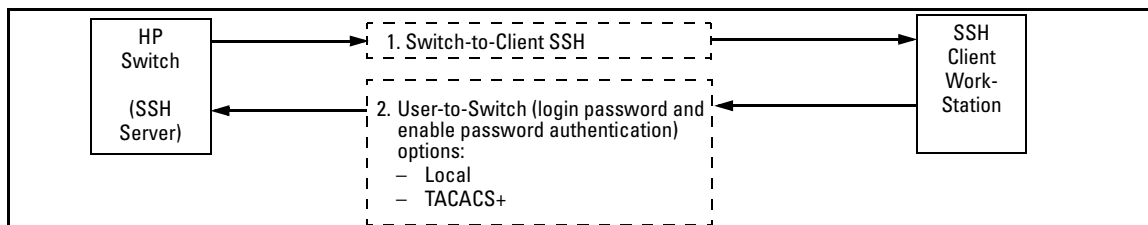
Figure 6-2.   Switch/User Authentication

Terminology

- **SSH Server:** An HP switch with SSH enabled.

- **Key Pair:** A pair of keys generated by the switch or an SSH client application. Each pair includes a public key, that can be read by anyone and a private key 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 portability and efficiency. SSHv2 client public-keys are typically stored in the PEM format. See figure 6-3 for an example of PEM-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 (**generate ssh [dsa | 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 6-7 and "4. Enabling SSH on the Switch and Anticipating SSH Client Contact Behavior" on page 6-12.)
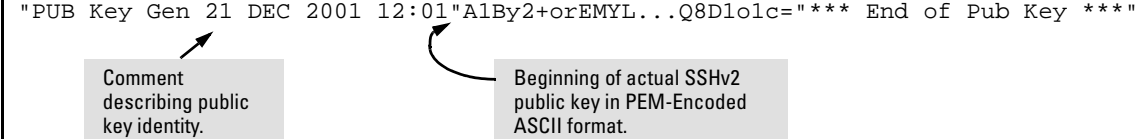
Prerequisite for Using SSH

Before using the switch as an SSH server, you must install a publicly or commercially available SSH client application on the computer(s) you use for management access to the switch. If you want client public-key authentication (page 6-1), then the client program must have the capability to generate or import keys.

Public Key Formats

Any client application you use for client public-key authentication with the switch must have the capability to export public keys. The switch can accept keys in the PEM-Encoded ASCII Format or in the Non-Encoded ASCII format.

```
"PUB Key Gen 21 DEC 2001 12:01"A1By2+orEMYL...Q8D1o1c="*** End of Pub Key ***"
```

Comment describing public key identity.

Beginning of actual SSHv2 public key in PEM-Encoded ASCII format.

Figure 6-3.   Example of Public Key in PEM-Encoded ASCII Format Common for SSHv2 Clients

Steps for Configuring and Using SSH for Switch and Client Authentication

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

Table 6-1. 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	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	none
	ssh enable tacacs	Yes	No	Yes	local or none
	ssh enable radius	Yes	No	Yes	local or none

[1] For **ssh login public-key**, 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 (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 6-20.

B. Switch Preparation

1. Assign a login (Operator) and enable (Manager) password on the switch (page 6-7).

2. Generate a public/private key pair on the switch (page 6-7).

   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 6-10).

4. Enable SSH on the switch (page 6-12).

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. If the primary option is local, the secondary option must be none.

      – Option B:

         Primary: Client public-key authentication (**login public-key** — page 6-19)
         Secondary: 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. If the primary option is local, the secondary option must be 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

■ Public keys generated on an SSH client must be exportable to the switch. The switch can only store 10 client key pairs.

■ 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.

■ 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*.

❧ With SSH running, the switch allows one console session and up to five other sessions (SSH and/or Telnet).

Configuring the Switch for SSH Operation

SSH-Related Commands in This Section	Page
show ip ssh	6-15
show crypto client-public-key [<manager	operator>] [keylist-str] [< babble	fingerprint>]	6-22
show crypto host-public-key [< babble	fingerprint >]	6-11
show authentication	6-18
crypto key < generate	zeroize > [cert [rsa] <keysize>	ssh [dsa	rsa [bits <keysize>]]	6-8
ip ssh	6-13
cipher <cipher-type>	6-13
filetransfer	6-14
mac	6-14
port < 1 - 65535|default >	6-13
timeout < 5 - 120 >	6-13
aaa authentication ssh	
login < local	tacacs	radius	public-key >	6-16, 6-17
< local	none >	6-16
enable < tacacs	radius	local >	6-16
< local	none >	6-16
copy tftp pub-key-file <tftp server IP> <public key file> [<append	manager	operator>]	6-22
clear crypto client-public-key [keylist-str]	6-23

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 >

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

Figure 6-4. 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 which should have access to the switch. Some SSH client applications automatically add 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.)

N o t e

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 (zeroing) 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**.)However, any active SSH sessions will continue to run, unless explicitly terminated with the CLI 'kill' command.

To Generate or Erase the Switch's Public/Private 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 <cert [rsa] <keysize> | ssh [dsa | rsa] bits <keysize>>

Installs authentication files for ssh or https server, or for autorun.

cert

Install RSA key for https certificate. See "Configuring the Switch for SSL Operation" on page 7-4 in this guide for more information.

ssh [dsa | rsa]

Install host key for ssh server. Specify the key type as DSA or RSA.

bits <keysize>

Specify the key size (in bits). See Table 6-2.

zeroize <ssh | cert | autorun [rsa]>

Erases the switch's public/private key pair and disables SSH operation.

show crypto host-public-key

Displays switch's public key. Displays the version 1 and version 2 views of the key.

[babble]

> *Displays hashes of the switch's public key in phonetic format. (See "Displaying the Public Key" on page 6-11.)*

[fingerprint]

> *Displays fingerprints of the switch's public key in hexadecimal format. (See "Displaying the Public Key" on page 6-11.)*

For example, to generate and display a new key:

```
HP Switch(config)# crypto key generate ssh rsa          Host Public
Installing new key pair.  If the key/entropy cache is   Key for the
depleted, this could take up to a minute.               Switch

HP Switch(config)# show crypto host-public-key

SSH host public key:

Ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAABAQDNwBMXZ9vYG+YxtV+KQeQQ+R8RKx47lxs14jPImBoV
qUmK1iyK3WGhv0Fsks9ZtATjidybwPjmfd2Lup1INUjemquL2C0YRXo4MULfrqioWTco4mlFliftypj5
gLvYeYFT/9CpoxMF4YrTUyjc90UfIf3et3fahaZ4KgLQjnXV1b8AWWsDZRv7niq8A1emqgts3HuAk59L
yFYsMkvLV2a7qoks1auqW8otRHgv/QsWtxfuoYUvVnXDiZFKeent34+sQnIvy1wJAmukCmM3zp434eTe
h87/2Oer612mlAwRq5qPFtTAZlsf8LGipjKftW6dq8uq9tYtTtQWO09Mi65L
```
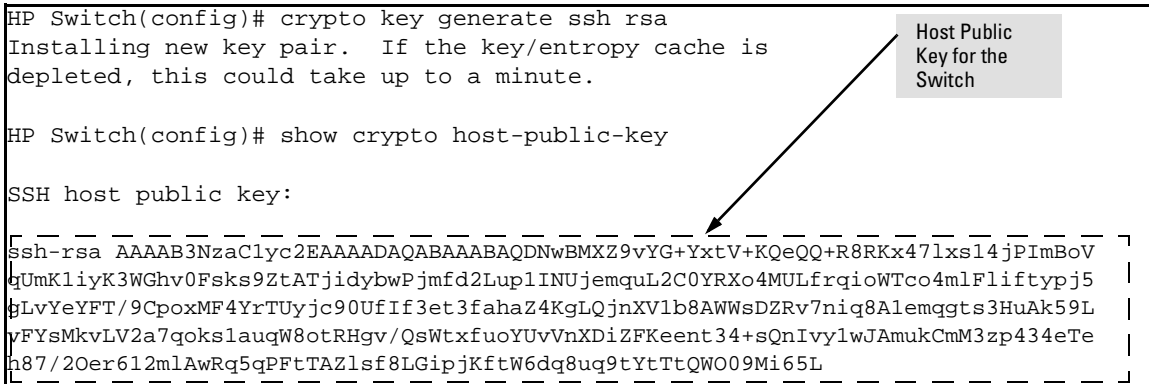
Figure 6-5.   Example of Generating a Public/Private Host Key Pair for the Switch

If you wish to compare the switch key to the key as stored in your client's known-hosts file, the formatting and comments need not match.

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.

Configuring Key Lengths

The **crypto key generate ssh** command allows you to specify the type and length of the generated host key. The size of the host key is platform-dependent as different switches have different amounts of processing power. The size is represented by the *<keysize>* parameter and has the values shown in Table 6-2. The default value is used if **keysize** is not specified.

Table 6-2. RSA/DSA Values for Various HP Switches

Platform	Maximum RSA Key Size (in bits)	DSA Key Size (in bits)
5400/3500/6200/8200/2910	1024, 2048, 3072 Default: 2048	1024
4200/2900/2810/2610/22520/510	1024, 2048 Default: 2048	1024
5300/2800/3400/2600	896	512

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.

The public key body follows the comments:

```
Comment          Public Key

ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAABAQCKzbhu89Lp5vmpuDwHmynUDnaBPPDtL9J8zcr4BLXi
k8G7VoMGghxJisLewYjFFUVVSOha8XCCGh4aVCur/HllP1enMGMGl/fu42oYDobmfgf/kJCK6U527rEU
gc+1VjxebjRGA5j0Bs8R2xo2k5ty6ZciMnDoEBeRcuoxd58dwzy7kU5hjQHSkubgwBBHYv7vmLtY0dPz
RHIRSoEezm5jFDhtF2EuubuB9adA43hggmKqciZZ8EfiOmKE9QwQCWLw2vUnXLhfFhxYwo7CoTir4yJA
lITSYFTrDd0G/pD67VfToz9DeHO163Yl9ukQQNXGES0LRK+8I1SZO33J7eb5
```

Figure 6-6.   Example of a Public Key Generated by the Switch

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 crypto host-public-key** command (figure 6-5).

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 changes or breaks in the text string. (A public key must be an unbroken ASCII string. Line breaks are not allowed Changes in the line breaks will corrupt the Key.) 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 6-7. Example of a Correctly Formatted Public Key

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:
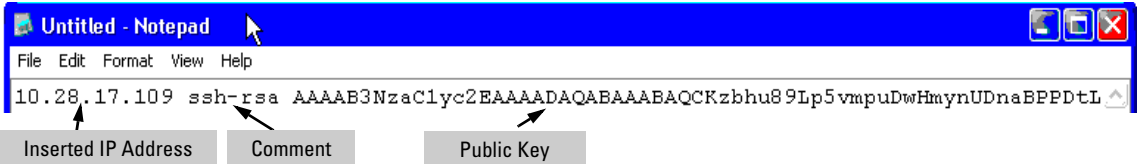
Figure 6-8. Example of a Switch Public Key Edited To Include the Switch's IP Address

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 length of the keys. (See figure 6-7 on page 6-11.)

■ **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 6-7 as follows:

```
HP Switch(config)# show crypto host-public-key babble    ← Phonetic "Hash" of Switch's Public Key
896 xozik-kohaf-daroh-fygas-byveb-bymiz-nupap-povaz-cesin-kafec-rixux
      host_sshl
896 xefes-hikot-kyher-cukuz-ba1ah-gezos-gumym-rezif-horib-cicyp-poxyx
      host_ssh2.pub
HP Switch(config)# show crypto host-public-key fingerprint       ← Hexadecimal
896 53:c0:14:75:72:84:90:cc:c8:ba:5e:ca:92:fc:c7:5c host_sshl    "Fingerprints" of
896 bf:fb:8a:d0:10:5a:48:57:61:f9:8a:6a:61:13:8a:fb host_ssh2.pub    the Same Switch
                                                                      Public Key
```

Figure 6-9. Examples of Visual Phonetic and Hexadecimal Conversions of the Switch's Public Key

The two commands shown in figure 6-9 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 has only one RSA host key. The 'babble' and 'fingerprint' options produce two hashes for the key--one that corresponds to the challenge hash you will see if connecting with a v1 client, and the other corresponding to the hash you will see if connecting with a v2 client. These hashes do not correspond to different keys, but differ only because of the way v1 and v2 clients compute the hash of the same RSA key. The switch always uses ASCII version (without 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.

N o t e

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 6-7.

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 configure SSH on the switch for client public-key authentication at the login (Operator) level. To enhance security, you should also configure local, TACACS+, or RADIUS authentication at the enable (Manager) level.

Refer to "5. Configuring the Switch for SSH Authentication" on page 6-15.

SSH Client Contact Behavior. At the first contact between the switch and an SSH client, if the switch's public key has not been copied into the client, then the client's first connection to the switch will question the connection and, for security reasons, provide the option of accepting or refusing. If it is safe to assume that an unauthorized device is not using the switch's IP address in an attempt to gain access to the client's data or network, the connection can be accepted. (As a more secure alternative, the client can be directly connected to the switch's serial port to download the switch's public key into the client. See the following Note.)

N o t e   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. This possibility can be removed 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 the client stores public keys, plus the knowledge of what key editing and file format might be required by the client application. However, if the 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 6-7.)

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 6-8).

Syntax: [no] ip ssh

Enables or disables SSH on the switch.

[cipher <cipher-type>]

Specify a cipher type to use for connection.

Valid types are:

- *aes128-cbc*
- *3des-cbc*
- *aes192-cbc*
- *aes256-cbc*
- *rijndael-cbc@lysator.liu.se*
- *aes128-ctr*
- *aes192-ctr*
- *aes256-ctr*

Default: All cipher types are available.

*Use the **no** form of the command to disable a cipher type.*

[filetransfer]

Enable/disable secure file transfer capability. SCP and SFTP secure file transfer will not function unless SSH is also enabled.

[mac <mac-type>]

Allows configuration of the set of MACs that can be selected. Valid types are:

- *hmac-md5*
- *hmac-sha1*
- *hmac-sha1-96*
- *hmac-md5-96*

Default: All MAC types are available.

*Use the **no** form of the command to disable a MAC type.*

[port < 1-65535 | default >]

*The TCP port number for SSH connections (default: 22). **Important:** See "Note on Port Number" on page 6-14.*

[timeout < 5 - 120 >]

Sets the maximum length of time (in seconds) allowed for initial protocol negotiation and authentication. Default: 120 seconds

Note on Port Number HP recommends using the default TCP 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 reserved TCP ports on the switch are 49, 80, 1506, and 1513.

```
HP Switch(config) ip ssh                 Enable SSH
HP Switch(config)# show ip ssh
  SSH Enabled    : Yes                Secure Copy Enabled : No
  TCP Port Number : 22                Timeout (sec)       : 120
  Host Key Type  : RSA               Host Key Size       : 1024

  Ciphers : aes128-cbc,3des-cbc,aes192-cbc,aes256-cbc,
           rijndael-cbc@lysator.liu.se,aes128-ctr,aes192-ctr,aes256-ctr
  MACs    : hmac-md5,hmac-sha1,hmac-sha1-96,hmac-md5-96

  Ses Type       | Source IP                                      Port
  --- --------  + ---------------------------------------------- -----
  1   console    |
  2   telnet     |                      With SSH running, the switch allows one console
  3   ssh        | 12.255.255.255       session and up to five other sessions (SSH and/or
  4   inactive   |                      Telnet). Web browser sessions are also allowed, but
  5   inactive   |                      do not appear in the show ip ssh listing.
  6   inactive   |
```

Figure 6-10. Example of Enabling IP SSH and Displaying the SSH Configuration

C a u t i o n   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, you should use SNMP version 3 only. If you need to increase the
security of your web interface see the section on SSL. Another security
measure is to use the Authorized IP Managers feature described in the switch's
Basic Operation 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 "Further Information on SSH
Client Public-Key Authentication" on page 6-19

N o t e

HP 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, ssh or serial port access enables full manager privileges.* See "1. Assigning a Local Login (Operator) and Enable (Manager) Password" on page 6-7.

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 | public-key >[< local | none | authorized >]

> *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**. *If the primary method is* **local**, *the secondary method must be* **none**.

> *The* **authorized** *option allows access without authentication.*

aaa authentication ssh enable < local | tacacs | radius | public-key>[< local | none | authorized >]

> *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**. *If the primary method is* **local**, *the secondary method must be* **none**.

> *The* **authorized** *option allows access without authentication.*

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 6-19.)

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 < *ipv4-address* | *ipv6-address* > < filename >

Copies a public key file into the switch.

aaa authentication ssh login public-key

Configures the switch to authenticate a client public-key at the login level with an optional secondary password method (default: **none***).*

Syntax: aaa authentication ssh enable < local | tacacs | radius | public-key > < local | none | authorized>

Configures a password method for the primary and second-ary enable (Manager) access. If you do not specify an optional secondary method, it defaults to **none**.

If the primary access method is **local**, *you can only specify* **none** *for a secondary access method.*

The **authorized** *option allows access without authentication.*

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 authenti-cation, with a Manager username of "1eader" and a password of "m0ns00n". To set up this operation you would configure the switch in a manner similar to the following:

```
                                    Configures Manager user- name and password.          Configures the switch
                                                                                          to allow SSH access
                                                                                          only for a client whose
                                                                                          public key matches
HP Switch(config)# password manager user-name leader                                      one of the keys in the
New password for Manager: ********                                                         public key file.
Please retype new password for Manager: ********                                          downloaded to the
HP Switch(config)# aaa authentication ssh login public-key none                           switch.
HP Switch(config)# aaa authentication ssh enable tacacs local
HP Switch(config)# coy tftp pub-key-file 10.33.18.117
HP Switch(config)# write memory                          Configures the primary and secondary
                                                         password methods for Manager (enable)
                    Copies a public key file named       access. (Becomes available after SSH
                    "Client-Keys.pub" into the switch.   access is granted to a client.)
```

Figure 6-11. Configuring for SSH Access Requiring a Client Public-Key Match and Manager Passwords

Figure 6-13 shows how to check the results of the above commands.

```
HP Switch(config)# show authentication

 Status and Counters - Authentication Information

  Login Attempts : 3
  Respect Privilege : Disabled

                  | Login       Login       Enable      Enable
  Access Task     | Primary     Secondary   Primary     Secondary
  ----------- + ---------- ---------- ---------- ----------
  Console         | Local       None        Local       None
  Telnet          | Local       None        Local       None
  Port-Access     | Local       None
  Webui           | Local       None        Local       None
  SSH             | PublicKey   None        Tacacs      Local
  Web-Auth        | ChapRadius None
  MAC-Auth        | ChapRadius None
```

Figure 6-12. Example of SSH Configuration with a Public Key

```
                                                    Shows the contents of the public key file
                                                    downloaded with the copy tftp command in
        Client Key Index Number                     figure 6-11. In this example, the file contains
                                                    two client public-keys.

HP Switch(config)# show crypto client-public-key
0,"Maden name [1024-hit rsa, Local_crypto @Localcrypto , Thu Nov 07 2002 2l:25:4
2]" ssh-rsa AAAAB3NzaClyc2EAAAADAQABAAAAgQCz9oNfgxMHUFEC6frSulSa4UhlEFznFhQqmgP2
9HXXp6NR/lQ0UmACtrFU+QDllEtM/YM9FrN/XvZH/kIxTdFC5exFX/Sl0tcRaFYzI9UjK80dBqVBKGKB
IyVEbCVrwIlqdAqhkaEX3dfTwIaPS2xArLCFIIsTZhnGVQTZDIJGABlfrlcrwI==
l,"[768-bit rsa, Local_Crypto @Localcrypto , Mon Dec 16 2002 23:0l:5l]" ssh-rsa
AAAAB3NzaClyc2EAAAADAQABAAAAYQD0tmzA32JBgeuFJN0iXI3hfooPKZ09JKCPQcXEVk7N+eKf9M0X
vnmfFuEpw/fpghlvsE66n8FDu7W/B2tKI/tgQLFgx7GiVcxNGhLiN0/pg5AuEym8EnclGu/LgAM9daM=
```

Figure 6-13. /Example of Client-Public-Key Listing

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 "RADIUS-Related Problems" in the Troubleshooting chapter of the *Management and Configuration Guide* for your switch.

Further Information on SSH Client Public-Key Authentication

The section titled "5. Configuring the Switch for SSH Authentication" on page 6-15 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 store up to ten public keys for authenticating clients. This requires storing an ASCII version of each client's public key (without 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 6-15.)

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

1. 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 a secure hash algorithm to create a hashinformation version.
   c. Returns the hash version to the switch.

7. The switch computes its own hash version of the data from 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.

3. Use **copy tftp** to copy the client-public-key file to the switch. The switch can hold 10 keys. The new key is appended to the client public-key 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 challenge-response authentication, and require an understanding of how to use your SSH client application.

```
ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAAABAQCKzbhu89Lp5vmpuDwNmynUDnaBPPDtL9J8zcr4BLXi
k8G7VoMGghxJisLewYjFFUVVSOha8XCCGh4aVCur/HllP1enMGMGl/fu42oYDobmfgf/kJCK6U527rEU
gc+1VjxebjRGA5j0Bs8R2xo2k5ty6ZciMnDoEBeRcuoxd58dwzy7kU5hjQHSkubgwBBHYv7vmLtY0dPz
RHIRSoEezm5jFDhtF2EuubuB9adA43hggmKqciZZ8EfiOmKE9QwQCWLw2vUnXLhfFhxYwo7CoTir4yJA
1ITSYFTrDd0G/pD67VfToz9DeHO163Yl9ukQQNXGES0LRK+8I1SZO33smith@support.cairns.com
```

Figure 6-14. Example of a Client Public Key

N o t e s
Comments in public key files, such as **smith@support.cairns.com** in figure 6-14, 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 6-14 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	See figure 6-7 on page 6-11. The key must be one unbroken 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. 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 or DSA	You can choose either RSA or DSA key types when using the **crypto key generate ssh** command. The **cert** and **autorun** parameters only use RSA key types.
Maximum Supported Public Key Length	3072 bits	Shorter key lengths allow faster operation, but also mean diminished security.
Maximum Host Key Sizes In Bits	RSA: 1024, 2048, 3072 DSA: 1024	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, placing a client-public-key into a Word for Windows text file and clicking on **File I Properties I Statistics**, lets you view the number of characters in the file, including spaces.

2. Copy the client's public key 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 authentica-

tion, 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. 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 or individually on a TFTP server to which the switch has access. 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@support.cairns.com** at the end of the key in figure 6-14 on page 6-21.)

Syntax: copy tftp pub-key-file <*ipv4-address* | *ipv6 -address*> <*filename*> [<append | manager | operator>]

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

*The **append** option adds the key(s) for operator access.*

*The **manager** option replaces the key(s) for manager access; follow with the 'append' option to add the key(s).*

*The **operator** option replaces the key(s) for operator access (default); follow with the 'append' option to add the key(s).*

show crypto client-public-key [<manager | operator>] [keylist-str] [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 phonetic hashes that are easier for visual comparisons.*

*The **fingerprint** option converts the key data to hexadecimal hashes that are for the same purpose.*

*The **keylist-str** selects keys to display (comma-delimited list).*

*The **manager** option allows you to select manager public keys*

*The **operator** option allows you to select operator public keys.*

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:

```
HP Switch(config)# copy tftp pub-key-file 10.38.252.195 clientkeys.txt

HP Switch(config)# show crypto client-public-key
0 "Maden name [1024-bit, rsa, Jamie_wilson@Jamiewilson , Thu Nov 07 2002 2l:25:4
2]  ssh-rsa AAAAB3NzaClYc2EAAAADAQABAAAgQCz9oNfqxMHUFEC6frSulSa4Uh1EFznFhQqmgP2
9HXYp6NR/1Q0UmACtrFU+QD1lEtM/YM9FrN/XvZH/kIxTdEc5exFX/Sl0tcRaFYzI9UjK80dBMqvBGKB
IyVEbCVwlqdAqbkaEX3df/WaPS2xArLCFHsTZhnCvQTZD0GABlfr1cw==
l,"[76S-bit rsa,  Jamie_wilson@Jamiewilson , Mon Dec 16 2002 23:0l:5l]" ssh-rsa
AAAAB3NzaClyc2EAAAADAQABAAAAYQD0tmzA32JBgeuFJN0iXI3bfooPKZ09JKCPQcXEVk7N+eKf9M0X
vnmfFuEpw/fpqh1vsE66n8FDu7W/B2tKH/tqQLFqx7GiVcxNHgLiN0/pq5AuEym8EnclGu/LgAM9daM=
```

Key Index Number

Figure 6-15. Example of Copying and Displaying a Client Public-Key File Containing Two Different Client Public Keys for the Same Client

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 remove the existing client public-key file or specific keys by executing the **clear crypto public-key** command.

Syntax: clear crypto public-key

Deletes the client-public-key file from the switch.

Syntax: clear crypto public-key 3

Deletes the entry with an index of 3 from the client-public-key file on the switch.

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 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.

Syntax: aaa authentication ssh login public-key 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.

C a u t i o n	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

Logging Messages

There are event log messages when a new key is generated and zeroized for the server:

ssh: New <num-bits> -bit [rsa | dsa] SSH host key installed
ssh: SSH host key zeroized

There are also messages that indicates when a client public key is installed or removed:

ssh: <num-bits>-bit [rsa | dsa] client public key [installed | removed] ([manager| operator] access) (key_comment)
Note: Only up to 39 characters of the key comment are included in the event log message.

Debug Logging

To add ssh messages to the debug log output, enter this command:

```
HP Switch# debug ssh LOGLEVEL
```

where LOGLEVEL is one of the following (in order of increasing verbosity):

- fatal
- error
- info
- verbose
- debug
- debug2
- debug3

Configuring Secure Socket Layer (SSL)

Overview

Feature	Default	Menu	CLI	Web
Generating a Self Signed Certificate on the switch	No	n/a	page 7-6	page 7-10
Generating a Certificate Request on the switch	No	n/a	n/a	page 7-12
Enabling SSL	Disabled	n/a	page 7-14	page 7-16

The switches covered in this guide use Secure Socket Layer Version 3 (SSLv3) and support for Transport Layer Security(TLSv1) to provide remote web access to the switches via encrypted paths between the switch and management station clients capable of SSL/TLS operation.

N o t e   HP Switches use SSL and TLS for all secure web transactions, and all references to SSL mean using one of these algorithms unless otherwise noted

SSL provides all the web functions but, unlike standard web access, SSL provides encrypted, authenticated transactions. The authentication type includes server certificate authentication with user password authentication.

N o t e   SSL in the switches covered in this guide is based on the OpenSSL software toolkit. For more information on OpenSSL, visit **www.openssl.com**.

Server Certificate authentication with User Password Authentication .   This option is a subset of full certificate authentication of the user and host. It occurs only if the switch has SSL enabled. As in figure 7-1, the switch authenticates itself to SSL enabled web browser. Users on SSL browser then authenticate themselves to the switch (operator and/or manger levels) by providing passwords stored locally on the switch or on a TACACS+ or RADIUS server. However, the client does not use a certificate to authenticate itself to the switch.
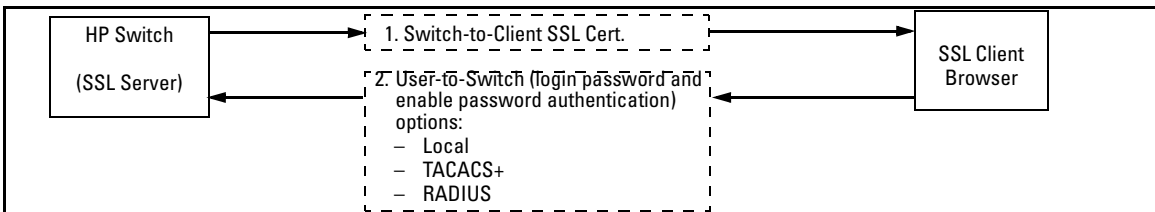
Figure 7-1.   Switch/User Authentication

SSL on the switches covered in this guide supports these data encryption methods:

■ 3DES (168-bit, 112 Effective)

■ DES (56-bit)

■ RC4 (40-bit, 128-bit)

N o t e :   HP Switches use RSA public key algorithms and Diffie-Hellman, and all references to a key mean keys generated using these algorithms unless otherwise noted

Terminology

■ **SSL Server:** An HP switch with SSL enabled.

■ **Key Pair:** Public/private pair of RSA keys generated by switch, of which public portion makes up part of server host certificate and private portion is stored in switch flash (not user accessible).

■ **Digital Certificate:** A certificate is an electronic "passport" that is used to establish the credentials of the subject to which the certificate was issued. Information contained within the certificate includes: name of the subject, serial number, date of validity, subject's public key, and the digital signature of the authority who issued the certificate. Certificates on HP switches conform to the X.509v3 standard, which defines the format of the certificate.

■ **Self-Signed Certificate:** A certificate not verified by a third-party certificate authority (CA). Self-signed certificates provide a reduced level of security compared to a CA-signed certificate.

■ **CA-Signed Certificate:** A certificate verified by a third party certificate authority (CA). Authenticity of CA-Signed certificates can be verified by an audit trail leading to a trusted root certificate.

■ **Root Certificate:** A trusted certificate used by certificate authorities to sign certificates (CA-Signed Certificates) and used later on to verify that authenticity of those signed certificates. Trusted certificates are distributed as an integral part of most popular web clients. (see browser documentation for which root certificates are pre-installed).

■ **Manager Level:** Manager privileges on the switch.

■ **Operator Level:** Operator privileges on the switch.

■ **Local password or username:** A Manager-level or Operator-level password configured in the switch.

■ **SSL Enabled:** (1)A certificate key pair has been generated on the switch (web interface or CLI command: **crypto key generate cert [key size]** (2) A certificate been generated on the switch (web interface or CLI command: **crypto host-cert generate self-signed [arg-list]**) and (3) SSL is enabled (web interface or CLI command: **web-management ssl**). (You can generate a certificate without enabling SSL, but you cannot enable SSL without first generating a Certificate.

Prerequisite for Using SSL

Before using the switch as an SSL server, you must install a publicly or commercially available SSL enabled web browser application on the computer(s) you use for management access to the switch.

Steps for Configuring and Using SSL for Switch and Client Authentication

The general steps for configuring SSL include:

A. Client Preparation

1. Install an SSL capable browser application on a management station you want to use for access to the switch. (Refer to the documentation provided with your browser.)

N o t e :   The latest versions of Microsoft Internet Explorer and Netscape web browser support SSL and TLS functionality. See browser documentation for additional details

B. Switch Preparation

1. Assign a login (Operator) and enable (Manager) password on the switch. (page 7-5)
2. Generate a host certificate on the switch. (page 7-6)
    i.   Generate certificate key pair
    ii. Generate host certificate

You need to do this only once. The switch's own public/private certificate key pair and certificate are stored in the switch's flash memory and are not affected by reboots or the erase startup-config command. (You can remove or replace this certificate, if necessary.) The certificate key pair and the SSH key pair are independent of each other, which means a switch can have two keys pairs stored in flash.

3. Enable SSL on the switch. (page 7-14)

4. Use your SSL enabled browser to access the switch using the switch's IP address or DNS name (if allowed by your browser). Refer to the documentation provided with the browser application.

General Operating Rules and Notes

■ Once you generate a certificate on the switch you should avoid re-generating the certificate without a compelling reason. Otherwise, you will have to re-introduce the switch's certificate on all management stations (clients) you previously set up for SSL access to the switch. In some situations this can temporarily allow security breaches.

■ The switch's own public/private certificate key pair and certificate are stored in the switch's flash memory and are not affected by reboots or the erase startup-config command

■ The public/private certificate key pair is not be confused with the SSH public/private key pair. The certificate key pair and the SSH key pair are independent of each other, which means a switch can have two keys pairs stored in flash

Configuring the Switch for SSL Operation

SSL-Related CLI Commands in This Section	Page
web-management ssl	page 7-16
show config	page 7-16
show crypto host-cert	7-9

SSL-Related CLI Commands in This Section	Page
crypto key	
generate cert [rsa] bits <1024 I 2048>	7-7
zeroize cert	7-7
crypto host-cert	
generate self-signed [arg-list]	7-7
zeroize	7-7

1. Assigning a Local Login (Operator) and Enabling (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.

Using the web browser interface To Configure Local Passwords. You can configure both the Operator and Manager password on one screen. To access the web browser interface, refer to the chapter titled "Using the Web Browser Interface" in the *Basic Operation Guide* for your switch.
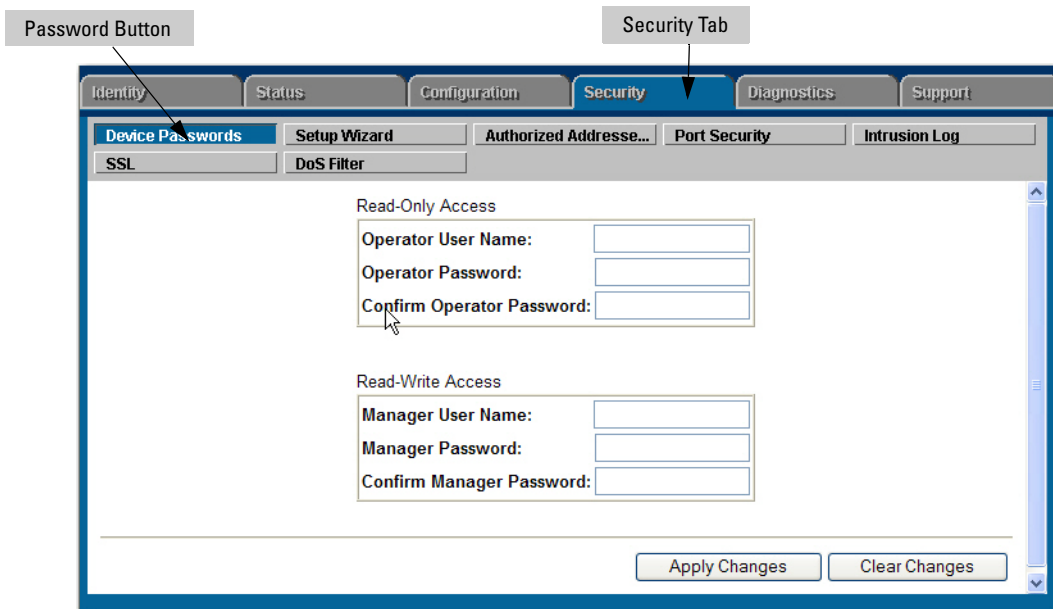
Figure 7-2. Example of Configuring Local Passwords

1. Proceed to the security tab and select device passwords button.

2. Click in the appropriate box in the Device Passwords window and enter user names and passwords. You will be required to repeat the password strings in the confirmation boxes.

   Both the user names and passwords can be up to 16 printable ASCII characters.

3. Click on **[Apply Changes]** button to activate the user names and passwords.

2. Generating the Switch's Server Host Certificate

You must generate a server certificate on the switch before enabling SSL. The switch uses this server certificate, along with a dynamically generated session key pair to negotiate an encryption method and session with a browser trying to connect via SSL to the switch. (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.)

The server certificate is stored in the switch's flash memory. The server certificate should be added to your certificate folder on the SSL clients who you want to have access to the switch. Most browser applications automatically add the switch's host certificate to there certificate folder on the first use. This method does allow for a security breach on the first access to the switch. (Refer to the documentation for your browser application.)

There are two types of certificated that can be used for the switch's host certificate. The first type is a self-signed certificate, which is generated and digitally signed by the switch. Since self-signed certificates are not signed by a third-party certificate authority, there is no audit trail to a root CA certificate and no fool-proof means of verifying authenticity of certificate. The second type is a certificate authority-signed certificate, which is digitally signed by a certificate authority, has an audit trail to a root CA certificate, and can be verified unequivocally

N o t e   There is usually a fee associated with receiving a verified certificate and the valid dates are limited by the root certificate authority issuing the certificate.

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

Removing (zeroizing) the switch's certificate key pair or certificate render the switch unable to engage in SSL operation and automatically disables SSL on the switch. (To verify whether SSL is enabled, execute **show config**.)

To Generate or Erase the Switch's Server Certificate with the CLI

Because the host certificate is stored in flash instead of the running-config file, it is not necessary to use **write memory** to save the certificate. Erasing the host certificate automatically disables SSL.

CLI commands used to generate a Server Host Certificate.

Syntax: crypto key generate cert [rsa] bits < 1024 | 2048 >

> *Generates a key pair for use in the certificate.*

crypto key zeroize cert

> *Erases the switch's certificate key and disables SSL operation.*

crypto host-cert generate self-signed [arg-list]

> *Generates a self signed host certificate for the switch. If a switch certificate already exists, replaces it with a new certificate. (See the Note, above.)*

crypto host-cert zeroize

> *Erases the switch's host certificate and disables SSL operation.*

To generate a host certificate from the CLI:
   i.   Generate a certificate key pair. This is done with the **crypto key generate cert** command. The default key size is 1024.

N o t e   If a certificate key pair is already present in the switch, it is not necessary to generate a new key pair when generating a new certificate. The existing key pair may be re-used and the crypto key generate cert command does not have to be executed.

   ii.   Generate a new self-signed host certificate. This is done with the **crypto host-cert generate self-signed [*Arg-List]*** command.

N o t e   When generating a self-signed host certificate on the CLI if there is not certificate key generated this command will fail.

Comments on Certificate Fields.

There are a number arguments used in the generation of a server certificate. table 7-1, "Certificate Field Descriptions" describes these arguments.

Table 7-1.Certificate Field Descriptions

Field Name	Description
Valid Start Date	This should be the date you desire to begin using the SSL functionality.
Valid End Date	This can be any future date, however good security practices would suggest a valid duration of about one year between updates of passwords and keys.
Common name	This should be the IP address or domain name associated with the switch. Your web browser may warn you if this field does not match the URL entered into the web browser when accessing the switch
Organization	This is the name of the entity (e.g. company) where the switch is in service.
Organizational Unit	This is the name of the sub-entity (e.g. department) where the switch is in service.
City or location	This is the name of the city where switch is in service
State name	This is the name of the state or province where switch is in service
Country code	This is the ISO two-letter country-code where switch is in service

For example, to generate a key and a new host certificate:

```
HP Switch(config)# crypto key generate cert bits 1024          Generate New Key
Installing new RSA key. If the key/entropy cache is
depleted, this could take up to a minute.                      Generate New Certificate
HP Switch(config)# crypto host-cert generate self-
signed
Validity start date [01/01/1970]: 01/01/2002
Validity end date   [01/01/2003]: 01/01/2004
Common name    [10.255.255.255]: 10.255.255.255               Enter certificate Arguments
Organization    [Company Name]: Hewlett Packard
Organizational unit  [Dept Name]: HP
City or locat1on         [City]: Roseille
State name              [State]: Ca
```

Figure 7-3.   Example of Generating a Self-Signed Server Host certificate on the CLI for the Switch.

N o t e s   "Zeroizing" the switch's server host certificate or key automatically disables
SSL (sets **web-management ssl** to **No**). Thus, if you zeroize the server host
certificate or key and then generate a new key and server certificate, you must
also re-enable SSL with the web-management ssl command before the switch
can resume SSL operation.

CLI Command to view host certificates.

Syntax: show crypto host-cert

Displays switch's host certificate

To view the current host certificate from the CLI you use the **show crypto host-cert** command.

For example, to display the new server host certificate:

```
HP Switch(config)#show crypto host-cert            Show host certificate command
Version: 1 (0x0)
Serial Number: 0 (0x0)
Issuer: CN=10.255.255.255, L=Roseville, ST=Ca, C=US, O=Hewlett Packard OU=HP Network
Validity
    Not Before: Jan 1 00:00:00 2002 GMT
    Not After : Jan 1 23:59:59 2000 GMT
Subject: CN=10:255.255.255, L=Roseville, ST=Ca, C=US, O=Hewlett Packard OU=HP Network
Subject Public Key Info:
   Public Key Algorithm: rsaEncryption
   RSA Public Key: (1024 bit)
      Modulus (1024 bit):
          00:db:18:4b:ce:3e:7d:5a:90:d8:a5:50:d5:2a:e9:
          60:78:d1:35:82:e9:27:71:5d:45:8d:0a:b9:b4:55:
          65:c7:d1:1c:4e:30:5e:20:a6:2d:62:9c:4c:cd:40:
          a0:6a:0b:cb:1c:ce:90:1c:2c:ad:26:fc:0b:07:ae:
          db:11:65:d6:47
      Exponent: 35 (0x23]
Signature Algorithm:  md5WithRSHEncryption
   d6:d0:98:6b:b9:a5:54:96:d9:be:fa:b9:99:f9:d8:6f:94:42:
   30:ea:c4:1d:88:e6:7b:19:18:22:84:f6:8c:ea:46:d7:ab:42:
   26:48:77:0c:60:57:8c:33:bc:08:d8:f7:c6:1f:ef:15:b7:24:
   f3:fa:92:b1:1f:7d:9e:c1:fd:83

MD5 Fingerprint: C969 E196 49C3 4609 AFC6 BDE1 2087 00A7
SHA1 Fingerprint: 93C7 0753 F805 26DC 4E39 EAF2 9C18 174F 7A63 E3C5
```
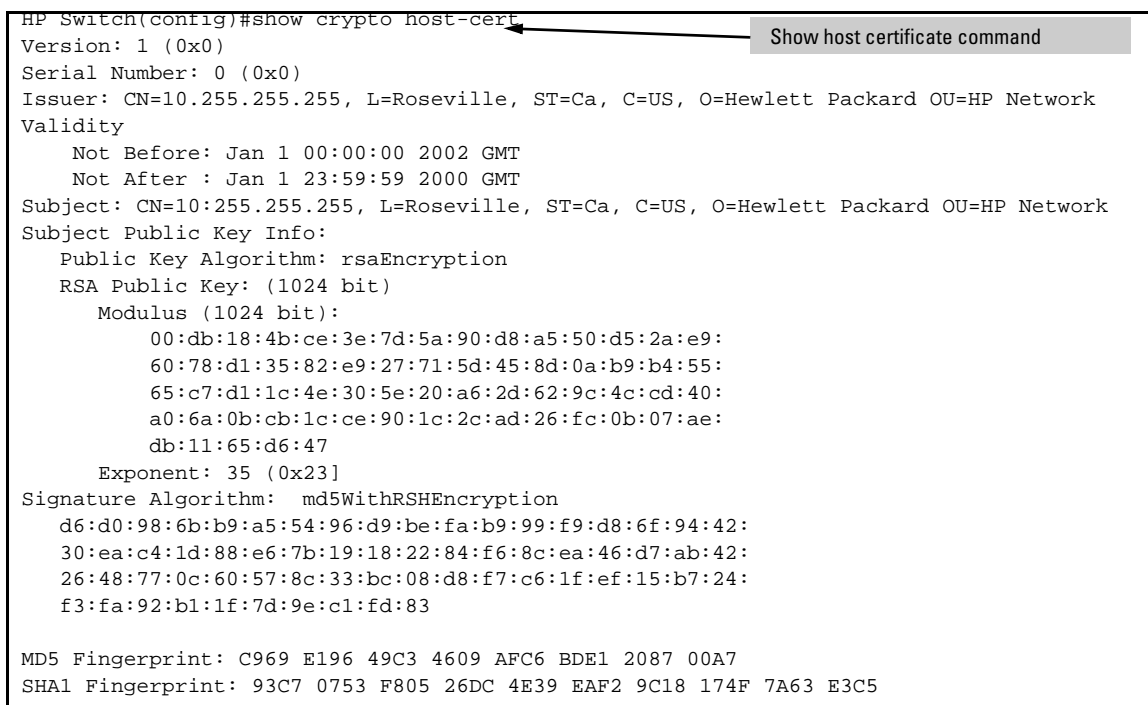
Figure 7-4.   Example of show crypto host-cert command

Generate a Self-Signed Host Certificate with the Web browser interface

You can configure SSL from the web browser interface. For more information on how to access the web browser interface refer to the chapter titled "Using the Web Browser Interface" in the *Basic Operation Guide* for your switch.

To generate a self signed host certificate from the web browser interface:

    i.   Proceed to the Security tab then the SSL button. The SSL configuration screen is split up into two halves. The left half is used in creating a new certificate key pair and (self-signed / CA-signed) certificate. The right half displays information on the currently installed certificate.

    ii.   Select the Generate Certificate button.

    iii.   Select Self signed certificate in the type box.

    iv.   Select the RSA key size desired. If you do not wish to generate a new key then just select current from the list.

    v.   Fill in remaining certificate arguments (refer to "To Generate or Erase the Switch's Server Certificate with the CLI" on page 7-7).

    vi.   Click on the **[Apply Changes]** button to generate a new certificate and key if selected.

N o t e   When generating a self-signed host certificate, if no key is present and the current option is selected in the RSA key size box and error will be generated. New key generation can take up to two minutes if the key queue is empty.

For example, to generate a new host certificate via the web browsers interface:
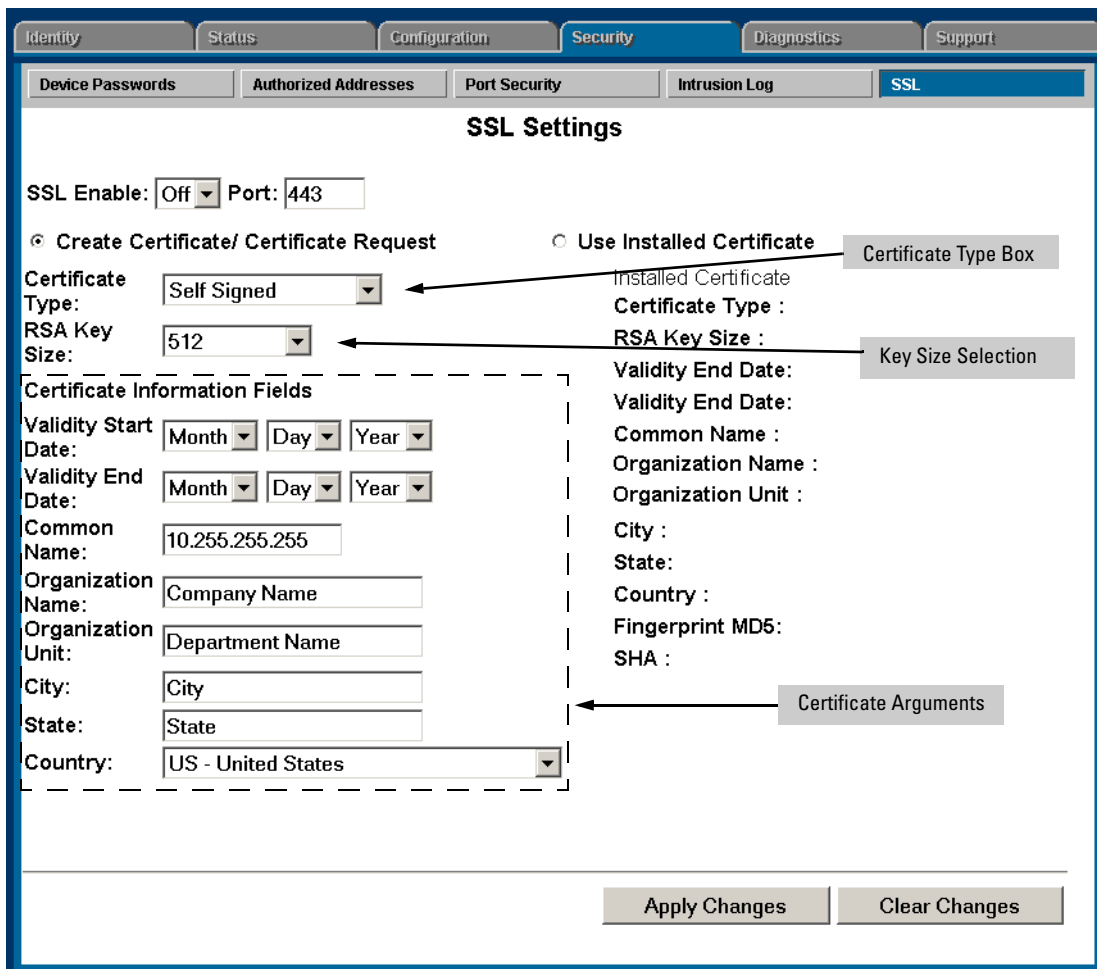
Figure 7-5. Self-Signed Certificate generation via SSL Web Browser Interface Screen

To view the current host certificate in the web browser interface:

1. Proceed to the **Security** tab

2. Then the **[SSL]** button
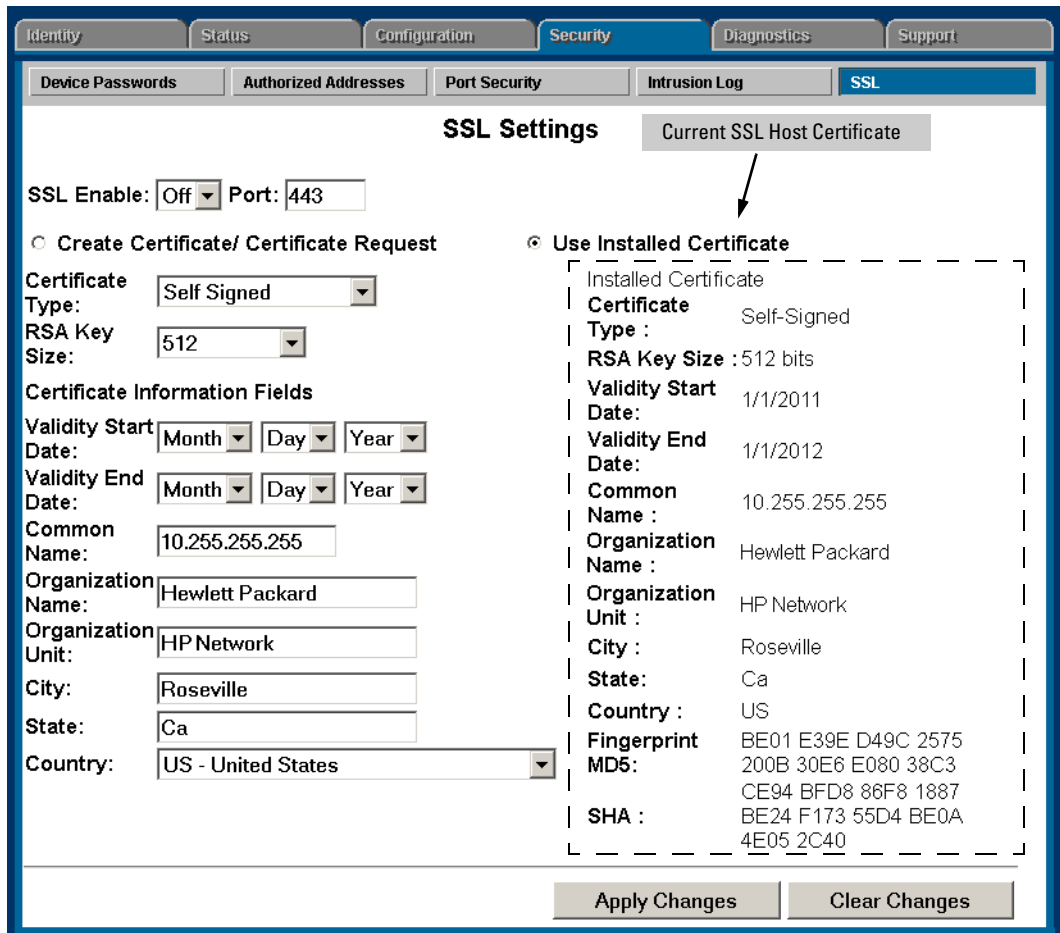
Figure 7-6. Web browser Interface showing current SSL Host Certificate

Generate a CA-Signed server host certificate with the Web browser interface

To install a CA-Signed server host certificate from the web browser interface. For more information on how to access the web browser interface, refer to the chapter titled "Using the Web Browser Interface" in the *Basic Operation Guide* for your switch.

The installation of a CA-signed certificate involves interaction with other entities and consists of three phases. The first phase is the creation of the CA certificate request, which is then copied off from the switch for submission to the certificate authority. The second phase is the actual submission process

that involves having the certificate authority verify the certificate request and then digitally signing the request to generate a certificate response (the usable server host certificate). The third phase is the download phase consisting of pasting to the switch web server the certificate response, which is then validated by the switch and put into use by enabling SSL

To generate a certificate request from the web browser interface:

    i.   Select the **Security** tab, then select the **[SSL]** button

    ii.   Select the **Create Certificate/Certificate Request** radio button.

    iii.   Select **Create CA Request** from the **Certificate Type** drop-down list.

    iv.   Select the key size from the RSA Key Size drop-down list. If you wish to re-use the current certificate key, select **Current** from the **RSA Key Size** drop-down list.

    v.   Fill in remaining certificate arguments (Refer to "Comments on Certificate Fields." on page 7-8.)

    vi.   Click on **[Apply Changes]** to create the certificate request. A new web browser page appears, consisting of two text boxes. The switch uses the upper text box for the certificate request text. The lower text box appears empty. You will use it for pasting in the certificate reply after you receive it from the certificate authority. (This authority must return a non- PEM encoded certificate request reply.

    vii.   After the certificate authority processes your request and sends you a certificate reply (that is, an installable certificate), copy and paste it into the lower text box.

    viii.   Click on the **[Apply Changes]** button to install the certificate.
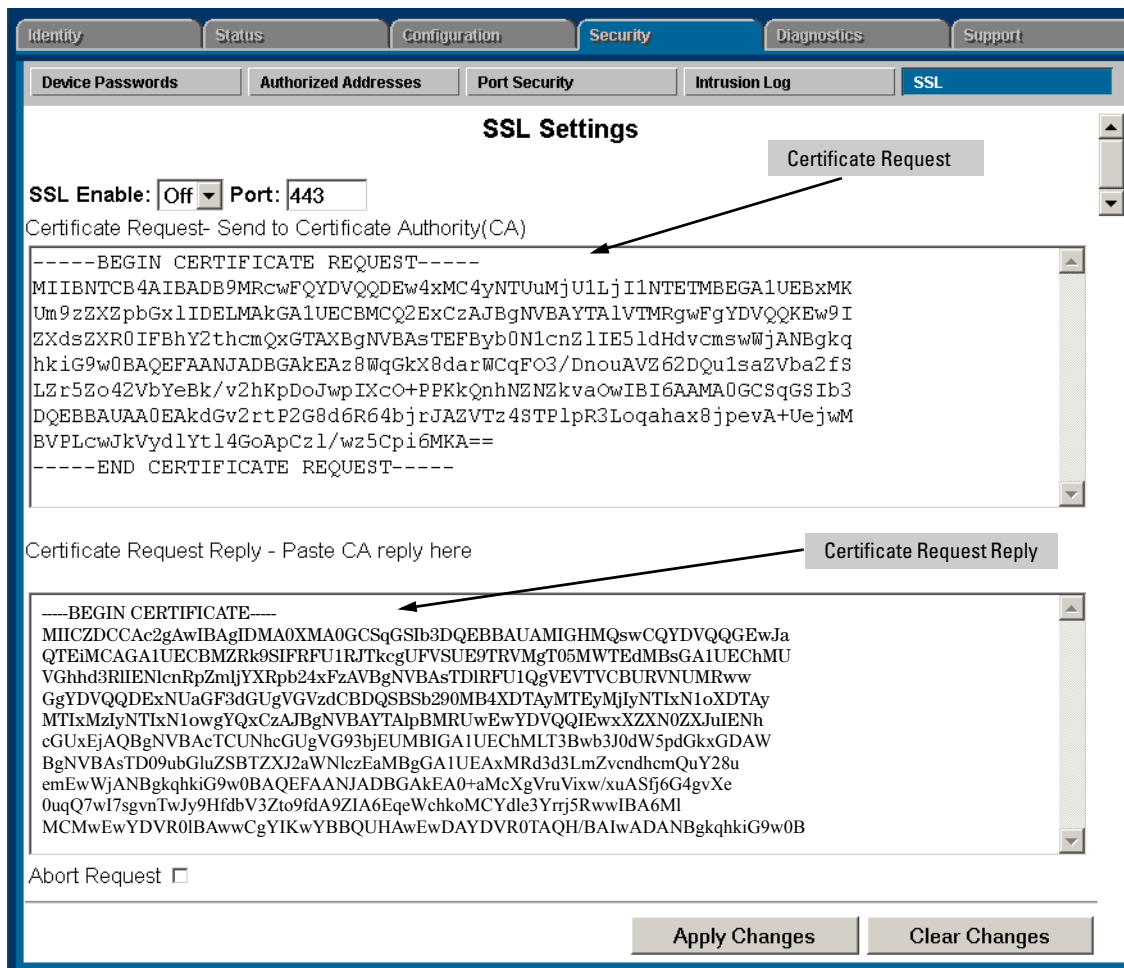
Figure 7-7. Request for Verified Host Certificate Web Browser Interface Screen

3. Enabling SSL on the Switch and Anticipating SSL Browser Contact Behavior

The **web-management ssl** command enables SSL on the switch and modifies parameters the switch uses for transactions with clients. After you enable SSL, the switch can authenticate itself to SSL enabled browsers. If you want to disable SSL on the switch, use the **no web-management ssl** command.

N o t e   Before enabling SSL on the switch you must generate the switch's host certificate and key. If you have not already done so, refer to "2. Generating the Switch's Server Host Certificate" on page 7-6.

When configured for SSL, the switch uses its host certificate to authenticate itself to SSL clients, however unless you disable the standard HP web browser interface with the **no web-management** command it will be still available for unsecured transactions.

SSL Client Contact Behavior. At the first contact between the switch and an SSL client, if you have not copied the switch's host certificate into the browser's certificate folder, your browser's first connection to the switch will question the connection and, for security reasons, give you the option of accepting or refusing. If a CA-signed certificate is used on the switch, for which a root certificate exists on the client browser side, then the browser will NOT prompt the user to ensure the validity of the certificate. The browser will be able to verify the certificate chain of the switch server certificate up to the root certificate installed in the browser, thus authenticating the switch unequivocally. 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.

N o t e   When an SSL 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. When using self-signed certificates with the switch, there is a possibility for a "man-in-the-middle" attack when connecting for the first time; that is, an unauthorized device could pose undetected as a switch, and learn the usernames and passwords controlling access to the switch. Use caution when connecting for the first time to a switch using self-signed certificates. Before accepting the certificate, closely verify the contents of the certificate (see browser documentation for additional information on viewing contents of certificate).

The security concern described above does not exist when using CA-signed certificates that have been generated by certificate authorities that the web browser already trusts

Using the CLI Interface to Enable SSL

Syntax: [no] web-management ssl

Enables or disables SSL on the switch.

[port < 1-65535 | default:443 >]

The TCP port number for SSL connections (default: 443). **Important:** *See "Note on Port Number" on page 7-17.*

show config

Shows status of the SSL server. When enabled **web-management ssl** *will be present in the config list.*

To enable SSL on the switch

1. Generate a Host certificate if you have not already done so. (Refer to "2. Generating the Switch's Server Host Certificate" on page 7-6.)

2. Execute the **web-management ssl** command.

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

■   Execute **no web-management ssl**.

■   Zeroize the switch's host certificate or certificate key. (page 7-7).

Using the Web Browser Interface to Enable SSL

To enable SSL on the switch
   i.   Proceed to the Security tab then the SSL button
   ii. Select SSL Enable to on and enter the TCP port you desire to connect on.
   iii. Click on the **[Apply Changes]** button to enable SSL on the port.

To disable SSL on the switch, do either of the following:
   i.   Proceed to the Security tab then the SSL button
   ii. Select SSL Enable to off .
   iii. Click on the **[Apply Changes]** button to enable SSL on the port.
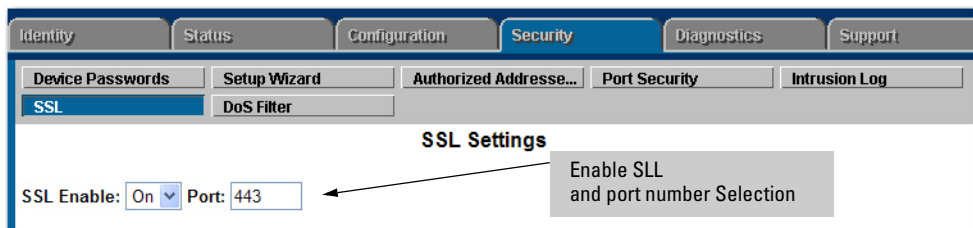
Figure 7-8.   Using the web browser interface to enable SSL and select TCP port number

Note on Port Number

HP recommends using the default IP port number (443). However, you can use **web-management ssl tcp-port** to specify any TCP port for SSL connections except those reserved for other purposes. Examples of reserved IP ports are 23 (Telnet) and 80 (http). Some other reserved TCP ports on the switches are 49, 80, 1506, and 1513.

Caution

SSL does not protect the switch from unauthorized access via the Telnet, SNMP, or the serial port. While 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 Telnet access (**no telnet**). If you need to increase SNMP security, use SNMP version 3 only for SNMP access. Another security measure is to use the Authorized IP Managers feature described in the switch's *Security 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.

Common Errors in SSL setup

Error During	Possible Cause
Generating host certificate on CLI	You have not generated a certificate key. (Refer to "CLI commands used to generate a Server Host Certificate" on page 7-7.)
Enabling SSL on the CLI or Web browser interface	You have not generated a host certificate. (Refer to "Generate a Self-Signed Host Certificate with the Web browser interface" on page 7-10.)
	You may be using a reserved TCP port. (Refer to "Note on Port Number" on page 7-17.)
Unable to Connect with SSL	You may not have SSL enabled (Refer to "3. Enabling SSL on the Switch and Anticipating SSL Browser Contact Behavior" on page 7-14.)
	Your browser may not support SSLv3 or TLSv1 or it may be disabled. (Refer to the documentation provided for your browser.)

8

IPv4 Access Control Lists (ACLs)

Introduction

An Access Control List (ACL) is a list of one or more Access Control Entries (ACEs) specifying the criteria the switch uses to either permit (forward) or deny (drop) IP packets traversing the switch's interfaces. This chapter describes how to configure, apply, and edit static IPv4 ACLs in a network populated with the switches covered by this guide, and how to monitor IPv4 ACL actions.

Note

This chapter describes static ACLs for IPv4 configuration and operation. Because the switches covered by this guide support IPv4/IPv6 dual-stack operation, simultaneous operation of statically configured IPv4 and IPv6 ACLs is supported in these switches. However:

■ IPv4 and IPv6 ACEs cannot be combined in the same static ACL.

■ IPv4 and IPv6 static ACLs do not filter each other's traffic.

In this chapter, unless otherwise noted:

■ The term "ACL" refers to static IPv4 ACLs.

■ Descriptions of ACL operation apply only to static IPv4 ACLs.

For information on configuring static IPv6 ACLs, refer to the chapter titled "IPv6 Access Control Lists (ACLs)" in the *IPv6 Configuration Guide* for your switch.

.

Feature	Default	CLI
Standard ACLs	None	8-39
Extended ACLs	None	8-48
Enable or Disable an ACL	n/a	8-62
Display ACL Data	n/a	8-76
Delete an ACL	n/a	8-64
Configure an ACL from a TFTP Server	n/a	8-86
Enable ACL Logging	n/a	8-91

IPv4 filtering with ACLs can help improve network performance and restrict network use by creating policies for:

- **Switch Management Access:** Permits or denies in-band management access. This includes limiting and/or preventing the use of designated protocols that run on top of IPv4, such as TCP, UDP, IGMP, ICMP, and others. Also included is control for application transactions based on source and destination IPv4 addresses and transport layer port numbers.

- **Application Access Security:** Eliminates unwanted traffic in a path by filtering IPv4 packets where they enter or leave the switch on specific VLAN interfaces.

IPv4 ACLs can filter traffic to or from a host, a group of hosts, or entire subnets.

Notes   IPv4 ACLs can enhance network security by blocking selected traffic, and can serve as part of your network security program. *However, because ACLs do not provide user or device authentication, or protection from malicious manipulation of data carried in IPv4 packet transmissions, they should not be relied upon for a complete security solution.*

IPv4 ACLs on the switches covered by this manual do not filter non-IPv4 traffic such as IPv6, AppleTalk, and IPX packets.

Overview of Options for Applying IPv4 ACLs on the Switch

To apply IPv4 ACL filtering, assign a configured IPv4 ACL to the interface on which you want traffic filtering to occur. VLAN and routed IPv4 traffic ACLs can be applied statically using the switch configuration.

Static ACLS

Static ACLs are configured on the switch. To apply a static ACL, you must assign it to an interface (VLAN or port).

VLAN ACL (VACL). A VACL is an ACL configured on a VLAN to filter traffic entering the switch on that VLAN interface and having a destination on the same VLAN.

Static Port ACL. A static port ACL is an ACL configured on a port to filter traffic entering the switch on that port, regardless of whether the traffic is routed, switched, or addressed to a destination on the switch itself.

Note   This chapter describes the IPv4 ACL applications you can statically configure on the switch. For information on static IPv6 ACL applications, refer to the chapter titled "IPv6 Access Control Lists (ACLs)" in the latest *IPv6 Configuration Guide* for your switch.

Command Summary for Standard IPv4 ACLs

Action	Command(s)	Page
Create a Standard, **Named** ACL *or* Add an ACE to the End of an Existing Standard, **Named** ACL	HP Switch(config)# ip access-list standard < *name-str* > HP Switch(config-std-nacl)# < deny | permit > < any | host <*SA* > | *SA*/< mask-length > | *SA* < *mask* >> [1] [log] [2]	8-41
Create a Standard, **Numbered** ACL or Add an ACE to the End of an Existing Standard, **Numbered** ACL	HP Switch(config)# access-list < 1-99 > < deny | permit > < any | host <*SA* > | *SA*/< mask-length > | *SA* < *mask* >> [log] [2]	8-44
Use a Sequence Number To Insert an ACE in a Standard ACL	HP Switch(config)# ip access-list standard < *name-str* | 1-99 > HP Switch(config-std-nacl)# 1-2147483647 < deny | permit > < any | host <*SA* > | *SA*/< mask-length > | *SA* < *mask* >> [1] [log] [2]	8-66
Use an ACE's Sequence Number To Delete the ACE from a Standard ACL	HP Switch(config)# ip access-list standard < *name-str* | 1-99 > HP Switch(config-std-nacl)# no < 1-2147483647 >	8-69
Resequence the ACEs in a Standard ACL	HP Switch(config)# ip access-list resequence < *name-str* | 1-99 > < 1-2147483647 > < 1-2147483646 >	8-70

Enter or Remove a Remark from a Standard ACL	HP Switch(config)# ip access-list standard < *name-str* | 1-99 >
	HP Switch(config-ext-nacl)# [remark < *remark-str* > | no < 1-2147483647 > remark]
	
	For numbered, standard ACLs only, the following **remark** *commands can be substituted for the above:*
	    HP Switch(config)# access-list < 1 - 99 > remark < remark-str >
	    HP Switch(config)# [no] access-list < 1 - 99 > remark
		8-71
		8-74

Delete a Standard ACL	HP Switch(config)# no ip access-list standard < *name-str* | 1-99 > *For numbered, standard ACLs, the following command can be substituted for the above:* HP Switch(config)# access-list < 1 - 99 > remark < remark-str >	8-64

[1] The mask can be in either dotted-decimal notation (such as 0.0.15.255) or CIDR notation (such as /20).
[2] The [log] function applies to both "deny" and "permit" ACLs, and generates a message when there is either a "deny" match or a "permit" match.

Command Summary for IPv4 Extended ACLs

Action	Command(s)	Page

Create an Extended, **Named** ACL
 	
or
 	
Add an ACE to the End of an Existing, Extended ACL
	HP Switch(config)# ip access-list extended < *name-str* l 100-199 >
	HP Switch(config-std-nacl)# < deny l permit >
	
	        < ip l *ip-protocol* l *ip-protocol-nbr* >
	           < any l host <*SA* > l *SA*/< mask-length > l *SA* < *mask* >> [1]
	           < any l host < *DA* > l *DA*/< mask-length > l *DA* < *mask* >> [1]
	        < tcp l udp >
	           < any l host <*SA* > l *SA*/< mask-length > l *SA* < *mask* >> [1]
	                [*comparison-operator* < *value* >]
	           < any l host <*DA* > l *DA*/< mask-length > l *DA* < *mask* >> [1]
	                [*comparison-operator* < *value* >]
	           [established]
	           [ack] [fin] [rst] [syn] [2]
	        < igmp >
	           < any l host <*SA* > l *SA*/< mask-length > l *SA* < *mask* >> [1]
	           < any l host < *DA* > l *DA*/< mask-length > l *DA* < *mask* >> [1]
	           [*igmp-packet-type*]
	        < icmp >
	           < any l host <*SA* > l *SA*/< mask-length > l *SA* < *mask* >> [1]
	           < any l host < *DA* > l *DA*/< mask-length > l *DA* < *mask* >> [1]
	           [[< 0 - 255 > [0 - 255]] l *icmp-message*]
	           [log] [3]
		8-50

Create an Extended, **Numbered** ACL or Add an ACE to the End of an Existing, **Numbered** ACL	HP Switch(config)# access-list < 100-199 > < deny l permit > < *ip-options* l*tcp/udp-options* l*igmp-options* l*icmp-options* > [log] [3] **Note:** *Uses the same IP, TCP/UDP, IGMP, and ICMP options as shown above for "Create an Extended, Named ACL".*	8-57
Insert an ACE by Assigning a Sequence Number	HP Switch(config)# ip access-list extended < *name-str* l 100-199 > HP Switch(config-ext-nacl)# 1-2147483647 < deny l permit > *Uses the options shown above for "Create an Extended, Named ACL".*	8-67
Delete an ACE by Specifying Its Sequence Number	HP Switch(config)# ip access-list extended < *name-str* l 100-199 > HP Switch(config-std-nacl)# no < 1-2147483647 >	8-69
Resequence the ACEs in an ACL	HP Switch(config)# ip access-list resequence < *name-str* l 100-199 > < 1-2147483647 > < 1-2147483646 >	8-70

[1] The mask can be in either dotted-decimal notation (such as 0.0.15.255) or CIDR notation (such as /20).
[2] TCP flag (control bit) options for destination TCP.
[3] The [log] function applies to both "deny" and "permit" ACLs, and generates a message when there is either a "deny" match or a "permit" match.

Action	Command(s)	Page
Enter or Remove a Remark	HP Switch(config)# ip access-list extended < *name-str*	100-199 > HP Switch(config-ext-nacl)# [remark < *remark-str* >	no remark] *For numbered, extended ACLs only, the following* **remark** *commands can be substituted for the above:* HP Switch(config)# access-list < 100 - 199 > remark < remark-str > HP Switch(config)# [no] access-list < 100 - 199 > remark	8-71 8-74
Delete an Extended ACL	HP Switch(config)# no ip access-list extended < *name-str*	100-199 > *For numbered, extended ACLs only, the following command can also be used:* HP Switch(config)# no access-list < 100 - 199 >	8-64

Command Summary for Enabling, Disabling, and Displaying ACLs

Enable or Disable a VACL	HP Switch(config)# [no] vlan < *vid* > ip access-group < *identifier* > < vlan >	
Enable or Disable a Static Port ACL	HP Switch((config)# [no] interface < *port-list*	Trk*x* > access-group < *identifier* > in HP Switch((eth-< *port-list* >	Trk*x* >)# [no] ip access-group < *identifier* > in	
Displaying ACL Data	HP Switch((config)# show access-list HP Switch((config)# show access-list < *acl-identifier* > [config] HP Switch((config)# show access-list config HP Switch((config)# show access-list ports < port-list	trk*x* > HP Switch((config)# show access-list vlan < *vid* >	8-76
Displaying or Clearing ACL Statistics	HP Switch(# < show	clear > statistics aclv4 < *acl-name-str* > port < port # > HP Switch(# < show	clear > statistics aclv4 < *acl-name-str* > vlan < *vid* > < in	out	vlan >	8-93

Terminology

Access Control Entry (ACE): A policy consisting of criteria and an action (permit or deny) to execute on a packet if it meets the criteria. The elements composing the criteria include:

- source IPv4 address and mask (standard and extended ACLs)
- destination IPv4 address and mask (extended ACLs only)
- either of the following:
  - all IPv4 traffic
  - IPv4 traffic of a specific IP protocol (extended ACLs only) (In the cases of TCP, UDP, ICMP, and IGMP, the criteria can include either all traffic of the protocol type or only the traffic of a specific sub-type within the protocol.)
- option to log packet matches with **deny** or **permit** ACEs

Access Control List (ACL): A list (or set) consisting of one or more explicitly configured Access Control Entries (ACEs) and terminating with an implicit "deny" ACE. ACL types include "standard" and "extended". See also "Standard ACL" and "Extended ACL". To filter IPv4 traffic, apply either type in any of the following ways:

- VACL: an ACL assigned to filter inbound traffic on a specific VLAN configured on the switch
- Static Port ACL: an ACL assigned to filter inbound traffic on a specific switch port

An ACL can be configured on a VLAN as a VACL, and on a port (or static trunk) as a static port ACL.

ACE: See "Access Control Entry".

ACL: See "Access Control List".

ACL ID: A number or alphanumeric string used to identify an ACL. A *standard* IPv4 ACL ID can have either an alphanumeric string or a number in the range of 1 to 99. An *extended* IPv4 ACL ID can have either an alphanumeric string or a number in the range of 100 to 199. See also "Identifier".

ACL Mask: Follows any IPv4 address (source or destination) listed in an ACE. Defines which bits in a packet's corresponding IPv4 addressing must exactly match the addressing in the ACE, and which bits need not match (wildcards). See also "How an ACE Uses a Mask To Screen Packets for Matches" on page 8-24.)

CIDR: This is the acronym for Classless Inter-Domain Routing.

DA: The acronym used in text to represent *Destination Address*. In an IPv4 packet, this is the destination address carried in the header, and identifies the destination intended by the packet's originator. In an extended ACE, this is the second of two addresses required by the ACE to determine whether there is a match between a packet and the ACE. See also "SA".

Deny: An ACE configured with this action causes the switch to drop a packet for which there is a match within an applicable ACL.

Extended ACL: This type of IPv4 Access Control List uses layer-3 IP criteria composed of source and destination addresses and (optionally) TCP/UDP port, ICMP, or IGMP criteria to determine whether there is a match with an IP packet.Extended ACLs require an alphanumeric name or an identification number (ID) in the range of 100 - 199.

IDENTIFIER*:** A term used in ACL syntax statements to represent either the name or number by which the ACL can be accessed. See also ***NAME-STR.

Implicit Deny: If the switch finds no matches between an IPv4 packet and the configured criteria in an applicable static ACL, then the switch denies (drops) the packet with an implicit **deny any** function (for standard ACLs) or an implicit **deny ip any any** function (for extended ACLs). You can preempt the Implicit Deny in a given ACL by configuring a **permit any** (standard) or **permit ip any any** (extended) as the last explicit ACE in the ACL. Doing so permits any IPv4 packet that is not explicitly permitted or denied by other ACEs configured sequentially earlier in the ACL. Unless otherwise noted, Implicit Deny refers to the "deny" function enforced by both standard and extended ACLs.

Inbound Traffic: For the purpose of defining where the switch applies IPv4 ACLs to filter traffic, inbound traffic is a packet that meets one of the following criteria:

- VLAN ACL (VACL): Inbound traffic is a packet entering the switch on a VLAN interface (or a subnet in a multinetted VLAN).

- Static Port ACL: Inbound traffic is a packet entering the switch on the port.

NAME-STR: A term used in extended ACL syntax statements to represent the "name string"; the alphanumeric string used to identify the ACL. See also **IDENTIFIER** and **ACL-ID**.

Named ACL: An ACL created with the **ip access-list < extended | standard > < name-str >** command and then populated using the **< deny | permit >** command in the Named ACL (**nacl**) CLI context. (Refer to "Entering the IPv4 "Named ACL" (nacl) Context" on page 8-41.)

Numbered ACL: An ACL created and initially populated by using the **access-list < extended < 100 - 199 > | standard > < 1-99 >>** command. (Refer to "Creating or Adding to an Standard, Numbered ACL" on page 8-45.) After a numbered ACL has been created, the switch manages it in the same way as a named ACL, meaning that it can be applied and edited in the same way as a named ACL.

Permit: An ACE configured with this action allows the switch to forward a packet for which there is a match within an applicable ACL.

Permit Any Forwarding: An ACE configured with this action causes the switch to forward IP packets that have not been permitted or denied by earlier ACEs in the list. In a standard ACL, this is **permit any**. In an extended ACL, it is **permit ip any any**.

remark-str: The term used in ACL syntax statements to represent the variable "remark string"; a set of alphanumeric characters you can include in a remark in an ACL. A remark string can include up to 100 characters and must be delimited by single or double quotes if any spaces are included in the string.

Rate-Limit Port ACLs (RL-PACLs): allows you to create an ACL and apply it on a per-port basis to rate-limit network traffic.

SA: The acronym for *Source Address*. In an IPv4 packet, this is the source IPv4 address carried in the IP header, and identifies the packet's sender. In a standard ACE, this is the IPv4 address used by the ACE to determine whether there is a match between a packet and the ACE. In an extended ACE, this is the first of two IPv4 addresses used by the ACE to determine whether there is a match between a packet and the ACE. See also "DA".

seq-#: The term used in ACL syntax statements to represent the sequence number variable used to insert an ACE within an existing list. The range allowed for sequence numbers is 1 - 2147483647.

Standard ACL: This type of access control list uses the layer-3 IP criteria of source IPv4 address to determine whether there is a match with an IPv4 packet. Standard ACLs require an alphanumeric name or an identification number (ID) in the range of 1- 99. See also *identifier* on page 8-8.

Static Port ACL: An ACL statically configured on a specific port, group of ports, or trunk. A static port ACL filters all incoming IPv4 traffic on the port, regardless of whether it is switched or routed.

VACL: See "VLAN ACL".

VLAN ACL (VACL): An ACL applied to all IPv4 traffic entering the switch on a given VLAN interface. See also "Access Control List".

Wildcard: The part of a mask that indicates the bits in a packet's IPv4 addressing that do not need to match the corresponding bits specified in an ACL. See also **ACL Mask** on page 8-8.

Overview

Types of IPv4 ACLs

A permit or deny policy for IPv4 traffic you want to filter can be based on source address alone, or on source address plus other factors.

Standard ACL: Use a standard ACL when you need to permit or deny IPv4 traffic based on source address only. Standard ACLs are also useful when you need to quickly control a performance problem by limiting IPv4 traffic from a subnet, group of devices, or a single device. (This can block all IPv4 traffic from the configured source, but does not hamper IPv4 traffic from other sources within the network.) A standard ACL uses an alphanumeric ID string or a numeric ID of 1 through 99. You can specify a single host, a finite group of hosts, or any host.

Extended ACL: Use an extended ACL when simple IPv4 source address restrictions do not provide the sufficient traffic selection criteria needed on an interface. Extended ACLs allow use of the following criteria:

■   source and destination IPv4 address combinations

■   IPv4 protocol options

Extended, named ACLs also offer an option to permit or deny IPv4 connections using TCP for applications such as Telnet, http, ftp, and others.

ACL Applications

ACL filtering is applied to IPv4 traffic as follows:

■   VLAN ACL (VACL): on a VLAN configured with a VACL, inbound IP traffic, regardless of whether it is switched or routed. On a multi-netted VLAN, this includes inbound IPv4 traffic from any subnet.

■   Static port ACL: any inbound IPv4 traffic on that port.

VACL Applications

VACLs filter any IPv4 traffic entering the switch on a VLAN configured with the "VLAN" ACL option.

> vlan < *vid* > ip access-group < *identifier* > vlan

For example, in figure 8-1, you would assign a VACL to VLAN 2 to filter all inbound switched or routed IPv4 traffic received from clients on the 10.28.20.0 network. In this instance, routed traffic received on VLAN 2 from VLANs 1 or 3 would not be filtered by the VACL on VLAN 2.

The subnet mask for this example is 255.255.255.0.

Configuring a VACL on VLAN 2 filters the inbound IPv4 traffic from clients B and, C for all switched and routed destinations on all VLANs on the switch. Traffic routed from VLANs 1 and 3 to VLAN 2 is not filtered by the VACL on VLAN 2 because the configured VACL applies only to IPv4 traffic entering the switch on VLAN 2 (and not from traffic routed from other VLANs configured on the switch.)

A 10.28.10.5

Switch with IPv4 Routing Enabled

VLAN 1
10.28.10.1
(One Subnet)

B 10.28.20.88

VLAN 2 with VACL
(One Subnet)
10.28.20.1

C 10.28.20.99

VLAN 3
(Multiple Subnets)
10.28.40.1   10.28.30.1

D 10.28.40.22

E 10.28.30.33

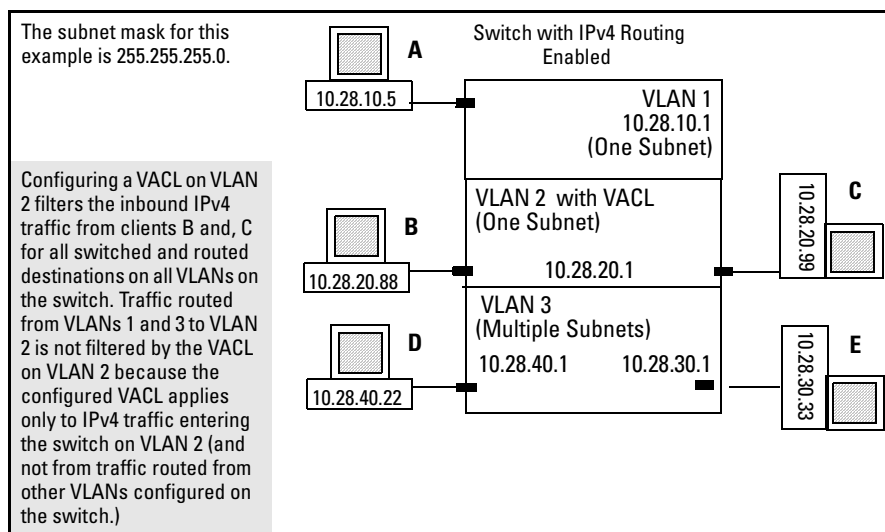**Figure 8-1. Example of VACL Filter Application to IPv4 Traffic Entering the Switch**

Note   The switch allows one VACL assignment configured per VLAN. This is in addition to any other ACL applications assigned to the VLAN or to ports in the VLAN.

Static Port ACL Applications

An IPv4 static port ACL filters any IPv4 traffic inbound on the designated port, Multiple ACLs on an Interface

The switch allows multiple ACL applications on an interface (subject to internal resource availability). This means that a port belonging to a given VLAN "X" can simultaneously be subject to all of the following:

Table 8-1.   Per-Interface Multiple ACL Assignments

ACL Type	ACL Application
IPv6 Static ACLs:	One static VACL for IPv6 traffic for VLAN "*X*" entering the switch through the port.
	One static port ACL for IPv6 traffic entering the switch on the port.
IPv4 Static ACLs:	one static VACL for IPv4 traffic for VLAN "*X*" entering the switch through the port
	one static port ACL for any IPv4 traffic entering the switch on the port

For a Packet To Be Permitted, It Must Have a Match with a "Permit" ACE in All Applicable ACLs Assigned to an Interface. On a given interface where multiple ACLs apply to the same traffic, a packet having a match with a **deny** ACE in any applicable ACL on the interface (including an implicit **deny any**) will be dropped.

For example, suppose the following is true:

■ Port A10 belongs to VLAN 100.

■ A static port ACL is configured on port A10.

■ A VACL is configured on VLAN 100.

An inbound, *switched* packet entering on port A10, with a destination on port A12, will be screened by the static port ACL and the VACL, regardless of a match with any **permit** or **deny** action. A match with a **deny** action (including an implicit deny) in either ACL will cause the switch to drop the packet. (If the packet has a match with explicit **deny** or **permit** ACEs in multiple ACLs and the **log** option is included in these ACEs, then a separate log event will occur for each match.)

Exception for Connection-Rate Filtering. Connection-rate filtering can be configured along with one or more other ACL applications on the same interface. In this case, a connection-rate match for a **filter** action is carried out according to the configured policy, regardless of whether any other ACLs on the interface have a match for a **deny** action. Also, if a connection-rate filter permits (**ignore** action) a packet, it can still be denied by another ACL on the interface.

Features Common to All ACL Applications

■ Any ACL can have multiple entries (ACEs).

■ You can apply any one ACL to multiple interfaces.

■ All ACEs in an ACL configured on the switch are automatically sequenced (numbered). For an existing ACL, entering an ACE without specifying a sequence number automatically places the ACE at the end of the list. Specifying a sequence number inserts the ACE into the list at the specified sequential location.

    • Automatic sequence numbering begins with "10" and increases in increments of 10. You can renumber the ACEs in an ACL and also change the sequence increment between ACEs.

    • The CLI **remark** command option allows you to enter a separate comment for each ACE.

■ A source or destination IPv4 address and a mask, together, can define a single host, a range of hosts, or all hosts.

■ Every ACL populated with one or more explicit ACEs includes an Implicit Deny as the last entry in the list. The switch applies this action to any packets that do not match other criteria in the ACL. (For standard ACLs, the Implicit Deny is **deny any**. For extended ACLs, it is **deny ip any any**.)

■ In any ACL, you can apply an ACL log function to ACEs that have an explicit "deny" or "permit" action. The logging occurs when there is a match on a "deny" or "permit" ACE. The switch sends ACL logging output to syslog, if configured, and, optionally, to a console session.

You can create ACLs for the switch configuration using either the CLI or a text editor. The text-editor method is recommended when you plan to create or modify an ACL that has more entries than you can easily enter or edit using the CLI alone. Refer to "Creating or Editing ACLs Offline" on page 8-86.

General Steps for Planning and Configuring ACLs

1. Identify the ACL action to apply. As part of this step, determine the best points at which to apply specific ACL controls. For example, you can improve network performance by filtering unwanted IPv4 traffic at the edge of the network instead of in the core. Also, on the switch itself, you can improve performance by filtering unwanted IPv4 traffic where it is inbound to the switch instead of outbound.

Traffic Source	ACL Application
IPv4 traffic entering the switch on a specific port	static port ACL (static-port assigned) for any inbound IPv4 traffic on a port from any source
switched or routed IPv4 traffic entering the switch on a specific VLAN	VACL (VLAN ACL)

2. Identify the traffic types to filter.
    *   The SA and/or the DA of traffic you want to permit or deny. This can be a single host, a group of hosts, a subnet, or all hosts.
    *   Traffic of a specific IPv4 protocol type (0-255)
    *   Any TCP traffic (only) for a specific TCP port or range of ports, including optional control of connection traffic based on whether the initial request should be allowed
    *   All UDP traffic or UDP traffic for a specific UDP port
    *   All ICMP traffic or ICMP traffic of a specific type and code
    *   All IGMP traffic or IGMP traffic of a specific type

3. Design the ACLs for the control points (interfaces) you have selected. Where you are using explicit "deny" ACEs, you can optionally use the VACL logging feature for notification that the switch is denying unwanted packets.

4. Configure the ACLs on the selected switches.

5. Assign the ACLs to the interfaces you want to filter, using the ACL application (static port ACL or VACL) appropriate for each assignment.

6. Test for desired results.

For more details on ACL planning considerations, refer to "Planning an ACL Application" on page 8-19.

IPv4 Static ACL Operation

Introduction

An ACL is a list of one or more Access Control Entries (ACEs), where each ACE consists of a matching criteria and an action (permit or deny). A static ACL applies only to the switch in which it is configured. ACLs operate on assigned interfaces, and offer these traffic filtering options:

■ IPv4 traffic inbound on a port.

■ IPv4 traffic inbound on a VLAN.

■ Routed IPv4 traffic entering or leaving the switch on a VLAN. (Note that ACLs do not screen traffic at the internal point where traffic moves between VLANs or subnets within the switch. Refer to "ACL Applications" on page 8-11.)

The following table lists the range of interface options:

Interface	ACL Application	Application Point	Filter Action
Port	Static Port ACL (switch configured)	inbound on the switch port	inbound IPv4 traffic
VLAN	VACL	entering the switch on the VLAN	inbound IPv4 traffic

Note   After you assign an IPv4 ACL to an interface, the default action on the interface is to implicitly deny IPv4 traffic that is not specifically permitted by the ACL. (This applies only in the direction of traffic flow filtered by the ACL.)

The Packet-filtering Process

Sequential Comparison and Action.   When an ACL filters a packet, it sequentially compares each ACE's filtering criteria to the corresponding data in the packet until it finds a match. The action indicated by the matching ACE (deny or permit) is then performed on the packet.

Implicit Deny.   If a packet does not have a match with the criteria in any of the ACEs in the ACL, the ACL denies (drops) the packet. If you need to override the implicit deny so that a packet that does not have a match will be

permitted, then you can use the "permit any" option as the last ACE in the ACL. This directs the ACL to permit (forward) packets that do not have a match with any earlier ACE listed in the ACL, and prevents these packets from being filtered by the implicit "deny any".

The ACL tries to apply the first ACE in the list. If there is not a match, it tries the second ACE, and so on. When a match is found, the ACL invokes the configured action for that entry (permit or drop the packet) and no further comparisons of the packet are made with the remaining ACEs in the list. This means that when an ACE whose criteria matches a packet is found, the action configured for that ACE is invoked, and any remaining ACEs in the ACL are ignored. *Because of this sequential processing, successfully implementing an ACL depends in part on configuring ACEs in the correct order for the overall policy you want the ACL to enforce.*
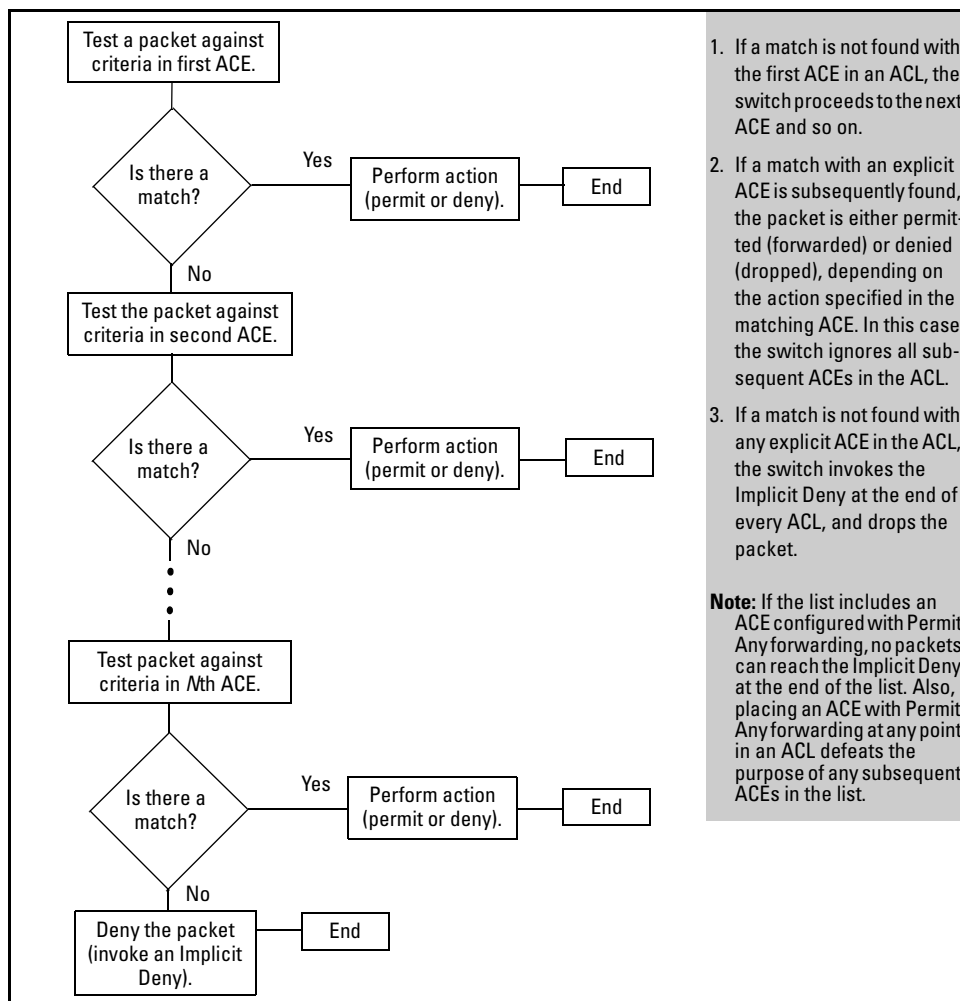
The flowchart shows the packet-filtering process:

Test a packet against criteria in first ACE. → Is there a match? → Yes → Perform action (permit or deny). → End

No ↓

Test the packet against criteria in second ACE. → Is there a match? → Yes → Perform action (permit or deny). → End

No ↓ · · ·

Test packet against criteria in *N*th ACE. → Is there a match? → Yes → Perform action (permit or deny). → End

No ↓

Deny the packet (invoke an Implicit Deny). → End

1. If a match is not found with the first ACE in an ACL, the switch proceeds to the next ACE and so on.

2. If a match with an explicit ACE is subsequently found, the packet is either permitted (forwarded) or denied (dropped), depending on the action specified in the matching ACE. In this case the switch ignores all subsequent ACEs in the ACL.

3. If a match is not found with any explicit ACE in the ACL, the switch invokes the Implicit Deny at the end of every ACL, and drops the packet.

Note: If the list includes an ACE configured with Permit Any forwarding, no packets can reach the Implicit Deny at the end of the list. Also, placing an ACE with Permit Any forwarding at any point in an ACL defeats the purpose of any subsequent ACEs in the list.

Figure 8-2. The Packet-Filtering Process in an ACL with *N* Entries (ACEs)

Note   The order in which an ACE occurs in an ACL is significant. For example, if an ACL contains six ACEs, but the first ACE allows Permit Any forwarding, then the ACL permits all IPv4 traffic, and the remaining ACEs in the list do not apply, even if they specify criteria that would make a match with any of the traffic permitted by the first ACE.

It is important to remember that all IPv4 ACLs configurable on the switch include an implicit **deny ip any**. That is, IPv4 packets that the ACL does not *explicitly* permit or deny will be *implicitly* denied, and therefore dropped instead of forwarded on the interface. If you want to preempt the implicit deny so that IPv4 packets not explicitly denied by other ACEs in the ACL will be permitted, insert an explicit "permit any" as the last ACE in the ACL. Doing so permits any packet not explicitly denied by earlier entries.

Planning an ACL Application

Before creating and implementing ACLs, you need to define the policies you want your ACLs to enforce, and understand how the ACL assignments will impact your network users.

Note   All IPv4 traffic entering the switch on a given interface is filtered by all ACLs configured for inbound traffic on that interface. For this reason, an inbound IPv4 packet will be denied (dropped) if it has a match with either an implicit or explicit **deny** in *any* of the inbound ACLs applied to the interface.

(See "An IPv4 static port ACL filters any IPv4 traffic inbound on the designated port, Multiple ACLs on an Interface" on page 8-12.)

IPv4 Traffic Management and Improved Network Performance

You can use ACLs to block traffic from individual hosts, workgroups, or subnets, and to block access to VLANs, subnets, devices, and services. Traffic criteria for ACLs include:

■   Switched and/or routed traffic

■   Any traffic of a specific IPv4 protocol type (0-255)

■   Any TCP traffic (only) for a specific TCP port or range of ports, including optional control of connection traffic based on whether the initial request should be allowed

■   Any UDP traffic or UDP traffic for a specific UDP port

■ Any ICMP traffic or ICMP traffic of a specific type and code

■ Any IGMP traffic or IGMP traffic of a specific type

Depending on the source and/or destination of a given IPv4 traffic type, you must also determine the ACL application(s) (VACL, or static port ACL) needed to filter the traffic on the applicable switch interfaces. Answering the following questions can help you to design and properly position IPv4 ACLs for optimum network usage.

■ What are the logical points for minimizing unwanted traffic, and what ACL application(s) should be used? In many cases it makes sense to prevent unwanted traffic from reaching the core of your network by configuring ACLs to drop the unwanted traffic at or close to the edge of the network. (The earlier in the network path you can block unwanted traffic, the greater the benefit for network performance.)

■ From where is the traffic coming? The source and destination of traffic you want to filter determines the ACL application to use (VACL or static port ACL).

■ What traffic should you explicitly block? Depending on your network size and the access requirements of individual hosts, this can involve creating a large number of ACEs in a given ACL (or a large number of ACLs), which increases the complexity of your solution.

■ What traffic can you implicitly block by taking advantage of the implicit **deny ip any** to deny traffic that you have not explicitly permitted? This can reduce the number of entries needed in an ACL.

■ What traffic should you permit? In some cases you will need to explicitly identify permitted traffic. In other cases, depending on your policies, you can insert an ACE with "permit any" forwarding at the end of an ACL. This means that all IPv4 traffic not specifically matched by earlier entries in the list will be permitted.

Security

ACLs can enhance security by blocking traffic carrying an unauthorized source IPv4 address (SA). This can include:

■ blocking access from specific devices or interfaces (port or VLAN)

■ blocking access to or from subnets in your network

■ blocking access to or from the internet

■ blocking access to sensitive data storage or restricted equipment

■ preventing specific IPv4, TCP, UDP, IGMP, and ICMP traffic types, including unauthorized access using functions such as Telnet, SSH, and web browser

You can also enhance switch management security by using ACLs to block IPv4 traffic that has the switch itself as the destination address (DA).

Caution   IPv4 ACLs can enhance network security by blocking selected traffic, and can serve as one aspect of maintaining network security. *However, because ACLs do not provide user or device authentication, or protection from malicious manipulation of data carried in IP packet transmissions, they should not be relied upon for a complete security solution.*

Note   Static IPv4 ACLs for the switches covered by this guide do not filter non-IPv4 traffic such as IPv6, AppleTalk, and IPX.

Guidelines for Planning the Structure of a Static ACL

After determining the filtering type (standard or extended) and ACL application (VACL, or static port ACL) to use at a particular point in your network, determine the order in which to apply individual ACEs to filter IPv4 traffic (For information on ACL applications, refer to "ACL Applications" on page 8-11.).

■ The sequence of ACEs is significant. When the switch uses an ACL to determine whether to permit or deny a packet on a particular VLAN, it compares the packet to the criteria specified in the individual Access Control Entries (ACEs) in the ACL, beginning with the first ACE in the list and proceeding sequentially until a match is found. When a match is found, the switch applies the indicated action (permit or deny) to the packet.

■ The first match in an ACL dictates the action on a packet. Subsequent matches in the same ACL are ignored. However, if a packet is permitted by one ACL assigned to an interface, but denied by another ACL assigned to the same interface, the packet will be denied on the interface.

■ On any ACL, the switch implicitly denies IPv4 packets that are not explicitly permitted or denied by the ACEs configured in the ACL. If you want the switch to forward a packet for which there is not a match in an ACL, append an ACE that enables Permit Any forwarding as the last ACE in the ACL. This ensures that no packets reach the Implicit Deny case for that ACL.

■ Generally, you should list ACEs from the most specific (individual hosts) to the most general (subnets or groups of subnets) unless doing so permits traffic that you want dropped. For example, an ACE allowing a small group of workstations to use a specialized printer should occur earlier in an ACL than an entry used to block widespread access to the same printer.

IPv4 ACL Configuration and Operating Rules

■ **VACLs and Switched or Routed IPv4 Traffic:** A VACL filters traffic entering the switch on the VLAN(s) to which it is assigned.

■ **Static Port ACLs:** A static port ACL filters traffic entering the switch on the port(s) or trunk(s) to which it is assigned.

■ **Per Switch ACL Limits for All ACL Types.** At a minimum an ACL must have one, explicit "permit" or "deny" Access Control Entry. You can configure up to 2048 IPv4 ACLs each for IPv4 and IPv6. The maximums are as follows:

   • Named (Extended or Standard) ACLs: Up to 2048 (minus any numeric standard or extended ACL assignments

   • Numeric Standard ACLs: Up to 99; numeric range: 1 - 99

   • Numeric Extended ACLs: Up to 100; numeric range: 100 - 199

   • The maximum number of ACEs supported by the switch is up to 3072 IPv4 ACEs, and up to 3072 IPv6 ACEs. The maximum number of ACEs allowed on a VLAN or port depends on the concurrent resource usage by multiple configured features. For more information, use the **show < qos | access-list > resources** command and/or refer to "Monitoring Shared Resources" on page 8-100. For a summary of IPv4 and IPv6 ACL resource limits, refer to the appendix covering scalability in the latest *Management and Configuration Guide* for your switch.

■ **Implicit Deny:** In any static IPv4 ACL, the switch automatically applies an implicit **deny ip any** that does not appear in **show** listings. This means that the ACL denies any IPv4 packet it encounters that does not have a match with an entry in the ACL. Thus, if you want an ACL to permit any packets that you have not expressly denied, you

must enter a **permit any** or **permit ip any any** as the last ACE in an ACL. Because, for a given packet the switch sequentially applies the ACEs in an ACL until it finds a match, any packet that reaches the **permit any** or **permit ip any any** entry will be permitted, and will not encounter the **deny ip any** ACE the switch automatically includes at the end of the ACL.

■ **Explicitly Permitting Any IPv4 Traffic:** Entering a **permit any** or a **permit ip any any** ACE in an ACL permits all IPv4 traffic not previously permitted or denied by that ACL. Any ACEs listed after that point do not have any effect.

■ **Explicitly Denying Any IPv4 Traffic:** Entering a **deny any** or a **deny ip any any** ACE in an ACL denies all IPv4 traffic not previously permitted or denied by that ACL. Any ACEs after that point have no effect.

■ **Replacing One ACL with Another Using the Same Application:** For a specific interface, the most recent ACL assignment using a given application replaces any previous ACL assignment using the same application on the same interface.

■ **Static Port ACLs:** These are applied per-port, per port-list, or per static trunk. Adding a port to a trunk applies the trunk's ACL configuration to the new member. If a port is configured with an ACL, the ACL must be removed before the port is added to the trunk. Also, removing a port from an ACL-configured trunk removes the ACL configuration from that port.

■ **VACLs:** These filter any IPv4 traffic entering the switch through any port belonging to the designated VLAN. VACLs do not filter traffic leaving the switch or being routed from another VLAN.

■ **VACLs Operate On Static VLANs:** You can assign an ACL to any VLAN that is statically configured on the switch. ACLs do not operate with dynamic VLANs.

■ **A VACL Affects All Physical Ports in a Static VLAN:** A VACL assigned to a VLAN applies to all physical ports on the switch belonging to that VLAN, including ports that have dynamically joined the VLAN.

How an ACE Uses a Mask To Screen Packets for Matches

When the switch applies an ACL to IPv4 traffic, each ACE in the ACL uses an IPv4 address and *ACL mask* to enforce a selection policy on the packets being screened. That is, the mask determines the range of IPv4 addresses (SA only or SA/DA) that constitute a match between the policy and a packet being screened.

What Is the Difference Between Network (or Subnet) Masks and the Masks Used with ACLs?

In common IPv4 addressing, a network (or subnet) mask defines which part of the address to use for the network number and which part to use for the hosts on the network. For example:

Address	Mask	Network Address	Host Address
10.38.252.195	255.255.255.0	first three octets	The fourth octet.
10.38.252.195	255.255.248.0	first two octets and the left-most five bits of the third octet	The right most three bits of the third octet and all bits in the fourth octet.

Thus, the bits set to 1 in a network mask define the part of an IPv4 address to use for the network number, and the bits set to 0 in the mask define the part of the address to use for the host number.

In an ACL, IPv4 addresses and masks provide criteria for determining whether to deny or permit a packet, or to pass it to the next ACE in the list. If there is a match, the configured deny or permit action occurs. If there is not a match, the packet is compared with the next ACE in the ACL. Thus, where a standard network mask defines how to identify the network and host numbers in an IPv4 address, the mask used with ACEs defines which bits in a packet's SA or DA must match the corresponding bits in the SA or DA listed in an ACE, and which bits can be *wildcards*.

Rules for Defining a Match Between a Packet and an Access Control Entry (ACE)

■   For a given ACE, when the switch compares an IPv4 address and corresponding mask in the ACE to an IPv4 address carried in a packet:

   •   **A mask-bit setting of 0 ("off")** requires that the corresponding bits in the packet's address and in the ACE's address must be the same. Thus, if a bit in the ACE's address is set to 1 ("on"), the same bit in the packet's address must also be 1.

   •   **A mask-bit setting of 1 ("on")** means the corresponding bits in the packet's address and in the ACE's address do not have to be the same. Thus, if a bit in the ACE's address is set to 1, the same bit in the packet's address can be either 1 or 0 ("on" or "off").

   For an example, refer to "Example of How the Mask Bit Settings Define a Match" on page 8-27.

■   In any ACE, a mask of all ones means *any* IPv4 address is a match. Conversely, a mask of all zeros means the *only* match is an IPv4 address identical to the host address specified in the ACE.

■   Depending on your network, a single ACE that allows a match with more than one source or destination IPv4 address may allow a match with multiple subnets. For example, in a network with a prefix of 31.30.240 and a subnet mask of 255.255.240.0 (the leftmost 20 bits), applying an ACL mask of 0.0.31.255 causes the subnet mask and the ACL mask to overlap one bit, which allows matches with hosts in two subnets: 31.30.224.0 and 31.30.240.0.

Bit Position in the Third Octet of Subnet Mask 255.255.240.0								
Bit Values	128	64	32	16	8	4	2	1
Subnet Mask Bits	1	1	1	1	n/a	n/a	n/a	n/a
Mask Bit Settings Affecting Subnet Addresses	0	0	0	**1 or 0**	n/a	n/a	n/a	n/a

This ACL supernetting technique can help to reduce the number of ACLs you need. You can apply it to a multinetted VLAN and to multiple VLANs. However, ensure that you exclude subnets that do not belong in the policy. If this creates a problem for your network, you can eliminate the unwanted match by making the ACEs in your ACL as specific as possible, and using multiple ACEs carefully ordered to eliminate unwanted matches.

■ Every IPv4 address and mask pair (source or destination) used in an ACE creates one of the following policies:

• **Any IPv4 address fits the matching criteria.** In this case, the switch automatically enters the address and mask in the ACE. For example:

access-list 1 deny any

produces this policy in an ACL listing:

Address	Mask
0.0.0.0	255.255.255.255

This policy states that every bit in every octet of a packet's SA is a wildcard, which covers any IPv4 address.

• **One IPv4 address fits the matching criteria.** In this case, you provide the address and the switch provides the mask. For example:

access-list 1 permit host 10.28.100.15

produces this policy in an ACL listing:

Address	Mask
10.28.100.15	0.0.0.0

This policy states that every bit in every octet of a packet's SA must be the same as the corresponding bit in the SA defined in the ACE.

• **A group of IPv4 addresses fits the matching criteria.** In this case you provide both the address and the mask. For example:

access-list 1 permit 10.28.32.1 0.0.0.31

Address	Mask
10.28.32.1	0.0.0.31

This policy states that:
   – In the first three octets of a packet's SA, every bit must be set the same as the corresponding bit in the SA defined in the ACE.
   – In the last octet of a packet's SA, the first three bits must be the same as in the ACE, but the last five bits are wildcards and can be any value.

■ Unlike subnet masks, the wildcard bits in an ACL mask need not be contiguous. For example, 0.0.7.31 is a valid ACL mask. However, a subnet mask of 255.255.248.224 is not a valid subnet mask.

Example of How the Mask Bit Settings Define a Match . Assume an ACE where the second octet of the mask for an SA is 7 (the rightmost three bits are "on", or "1") and the second octet of the corresponding SA in the ACE is 31 (the rightmost five bits). In this case, a match occurs when the second octet of the SA in a packet being filtered has a value in the range of 24 to 31. Refer to table 8-2, below.

Table 8-2.   Example of How the Mask Defines a Match

Location of Octet	Bit Position in the Octet							
	128	64	32	16	8	4	2	1
SA in ACE	0	0	0	1	1	1	1	1
Mask for SA	0	0	0	0	0	1	1	1
Corresponding Octet of a Packet's SA	0	0	0	1	1	0/1	0/1	0/1

The shaded area indicates bits in the packet that must exactly match the bits in the source address in the ACE. Wherever the mask bits are ones (wildcards), the corresponding address bits in the packet can be any value, and where the mask bits are zeros, the corresponding address bits in the packet must be the same as those in the ACE. **Note:** This example covers only one octet of an IPv4 address. An actual ACE applies this method to all four octets of the address.

Example of Allowing Only One IPv4 Address ("Host" Option). Suppose, for example, that you have configured the ACL in figure 8-3 to filter inbound packets on VLAN 20. Because the mask is all zeros, the ACE policy dictates that a match occurs only when the source address on such packets is identical to the address configured in the ACE.
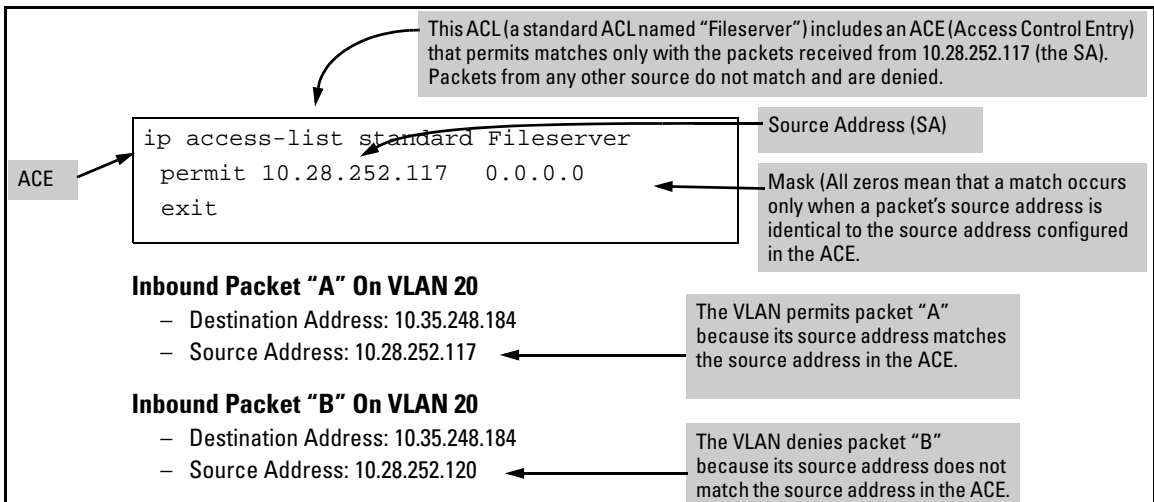
This ACL (a standard ACL named "Fileserver") includes an ACE (Access Control Entry) that permits matches only with the packets received from 10.28.252.117 (the SA). Packets from any other source do not match and are denied.

```
ip access-list standard Fileserver
  permit 10.28.252.117   0.0.0.0
  exit
```

ACE

Source Address (SA)

Mask (All zeros mean that a match occurs only when a packet's source address is identical to the source address configured in the ACE.

Inbound Packet "A" On VLAN 20
 – Destination Address: 10.35.248.184
 – Source Address: 10.28.252.117

The VLAN permits packet "A" because its source address matches the source address in the ACE.

Inbound Packet "B" On VLAN 20
 – Destination Address: 10.35.248.184
 – Source Address: 10.28.252.120

The VLAN denies packet "B" because its source address does not match the source address in the ACE.

Figure 8-3. Example of an ACL with an Access Control Entry (ACE) that Allows Only One Source Address

Examples Allowing Multiple IPv4 Addresses. Table 8-3 provides examples of how to apply masks to meet various filtering requirements.

Table 8-3.   Example of Using an IPv4 Address and Mask in an Access Control Entry

Address in the ACE	Mask	Policy for a Match Between a Packet and the ACE	Allowed Addresses
A: 10.38.252.195	0.0.0.255	Exact match in first three octets only.	10.38.252.< 0-255 > (See row A in table 8-4, below.)
B: 10.38.252.195	0.0.7.255	Exact match in the first two octets and the leftmost five bits (248) of the third octet.	10.38.< 248-255 >.< 0-255 > (In the third octet, only the rightmost three bits are wildcard bits. The leftmost five bits must be a match, and in the ACE, these bits are all set to 1. See row B in table 8-4, below.)
C: 10.38.252.195	0.0.0.0	Exact match in all octets.	10.38.252.195 (There are no wildcard bits in any of the octets. See row C in table 8-4, below.)
D: 10.38.252.195	0.15.255.255	Exact match in the first octet and the leftmost four bits of the second octet.	10.< 32-47 >.< 0-255 >.<0-255> (In the second octet, the rightmost four bits are wildcard bits. See row D in table 8-4, below.)

Table 8-4.   Mask Effect on Selected Octets of the IPv4 Addresses in Table 8-3

Addr	Octet	Mask	Octet Range	128	64	32	16	8	4	2	1
A	3	**0** all bits	252	1	1	1	1	1	1	0	0
B	3	**7** last 3 bits	248-255	1	1	1	1	1	0 or 1	0 or 1	0 or 1
C	4	**0** all bits	195	1	1	0	0	0	0	1	1
D	2	**15** last 4 bits	32-47	0	0	1	0	0 or 1	0 or 1	0 or 1	0 or 1

Shaded areas indicate bit settings that must be an exact match.

If there is a match between the policy in the ACE and the IPv4 address in a packet, then the packet is either permitted or denied, according to how the ACE is configured. If there is not a match, the next ACE in the ACL is then applied to the packet. The same operation applies to a destination IPv4 address (DA) used in an extended ACE. (Where an ACE includes both source and destination addresses, there is one address/ACL-mask pair for the source address, and another address/ACL-mask pair for the destination address. See "Configuring and Assigning an IPv4 ACL" on page 8-29.)

CIDR Notation. For information on using CIDR notation to specify ACL masks, refer to "Using CIDR Notation To Enter the IPv4 ACL Mask" on page 8-38.

Configuring and Assigning an IPv4 ACL

ACL Feature	Page
Configuring and Assigning a Standard ACL	8-39
Configuring and Assigning an Extended ACL	8-48
Enabling or Disabling ACL Filtering	8-62

Overview

General Steps for Implementing ACLs

1. Configure one or more ACLs. This creates and stores the ACL(s) in the switch configuration.

2. Assign an ACL. This step uses one of the following applications to assign the ACL to an interface:

   - VACL (any IPv4 traffic entering the switch on a given VLAN)
   - Static Port ACL (any IPv4 traffic entering the switch on a given port, port list, or static trunk)

Options for Permit/Deny Policies

The permit or deny policy for IPv4 traffic you want to filter can be based on source address alone, or on source address plus other IPv4 factors.

■　**Standard ACL:** Uses only a packet's source IPv4 address as a criterion for permitting or denying the packet. For a standard ACL ID, use either a unique numeric string in the range of 1-99 or a unique name string of up to 64 alphanumeric characters.

■　**Extended ACL:** Offers the following criteria as options for permitting or denying a packet:

•　source IPv4 address

•　destination IPv4 address

•　IPv4 protocol options:
　–　Any IPv4 traffic
　–　Any traffic of a specific IPv4 protocol type (0-255)
　–　Any TCP traffic (only) for a specific TCP port or range of ports, including optional use of TCP control bits or control of connection (established) traffic based on whether the initial request should be allowed
　–　Any UDP traffic (only) or UDP traffic for a specific UDP port
　–　Any ICMP traffic (only) or ICMP traffic of a specific type and code
　–　Any IGMP traffic (only) or IGMP traffic of a specific type

For an extended ACL ID, use either a unique number in the range of 100-199 or a unique name string of up to 64 alphanumeric characters.

Carefully plan ACL applications before configuring specific ACLs. For more on this topic, refer to "Planning an ACL Application" on page 8-19.

ACL Configuration Structure

After you enter an ACL command, you may want to inspect the resulting configuration. This is especially true where you are entering multiple ACEs into an ACL. Also, it is helpful to understand the configuration structure when using later sections in this chapter.

The basic ACL structure includes four elements:

1.　ACL identity and type: This identifies the ACL as **standard** or **extended** and shows the ACL name or number.

2.　Optional **remark** entries.

3. One or more deny/permit list entries (ACEs): One entry per line.

Element	Notes
Type	Standard or Extended
Identifier	• Alphanumeric; Up to 64 Characters, Including Spaces • Numeric: 1 - 99 (Standard) or 100 - 199 (Extended)
Remark	Allows up to 100 alphanumeric characters, including blank spaces. (If any spaces are used, the remark must be enclosed in a pair of single or double quotes.) A remark is associated with a particular ACE and will have the same sequence number as the ACE. (One remark is allowed per ACE.) Refer to "Attaching a Remark to an ACE" on page 8-71.
Maximum ACEs Per per Switch	The upper limit on ACEs supported by the switch depends on the concurrent resource usage by configured ACL, QoS, IDM, Mirroring, virus-throttling, and other features. Refer to "Monitoring Shared Resources" on page 8-100.

4. Implicit Deny: Where an ACL is in use, it denies any packets that do not have a match with the ACEs explicitly configured in the list. The Implicit Deny does not appear in ACL configuration listings, but always functions when the switch uses an ACL to filter packets. (You cannot delete the Implicit Deny, but you can supersede it with a **permit any** or **permit ip any any** statement.)

Standard ACL Structure

Individual ACEs in a standard ACL include only a permit/deny statement, the source addressing, and an optional **log** command (available with "deny" or "permit" statements).

```
ip access-list standard < identifier >"
  [ [ seq-# ] remark < remark-str >]
  < permit | deny > < SA > [log]
  .
  .
  .
  < Implicit Deny >
  exit
```

Note: The optional **log** function is available for either explicit "deny" or explicit "permit" ACEs.

Figure 8-4. Example of the General Structure for a Standard ACL

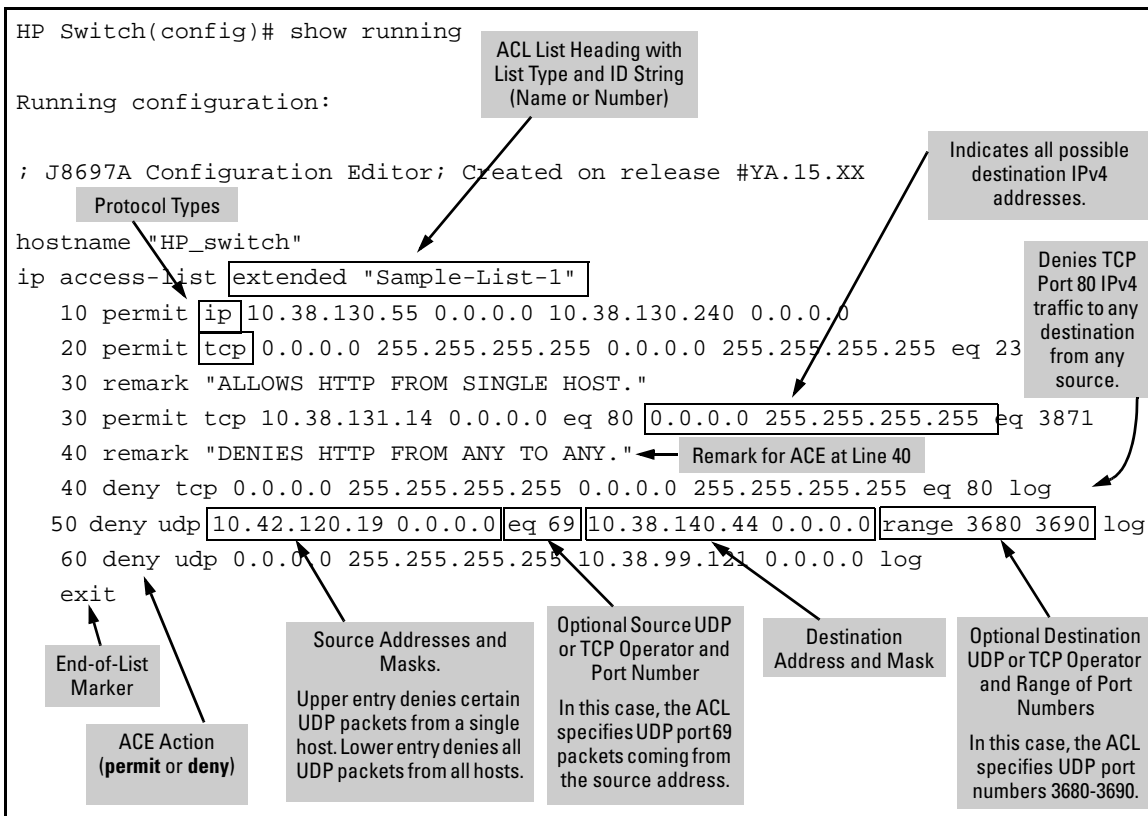For example, figure 8-5 shows how to interpret the entries in a standard ACL.

```
                    HP Switch(Config)# show running
                    .                        ACL List Heading with List Type and
                    .                           Identifier (Name or Number)
                    .
                    ip access-list standard "Sample-List"
  ACE Action        10 deny 10.28.150.77 0.0.0.0 log
(permit or deny)    20 permit 10.28.150.1 0.0.0.255
                    exit

End-of-List Marker          Source Address      Mask      Optional Logging
                                                          Command
```
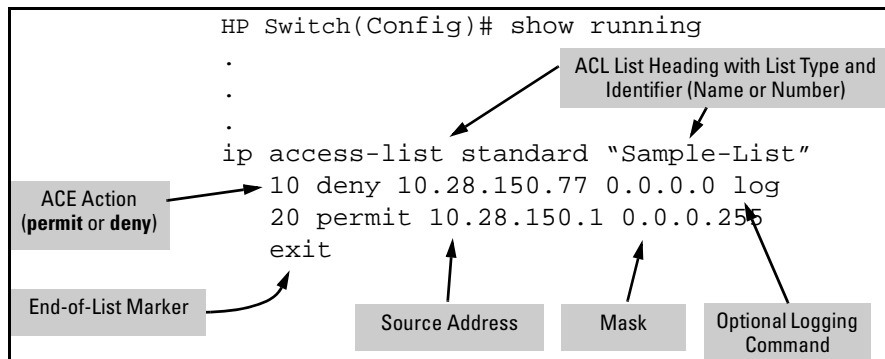
Figure 8-5. Example of a Displayed Standard ACL Configuration with Two ACEs

Extended ACL Configuration Structure

Individual ACEs in an extended ACL include:

- A permit/deny statement

- Source and destination IPv4 addressing

- Choice of IPv4 criteria

- Optional ACL **log** command (for **deny** or **permit** entries)

- Optional remark statements

```
ip access-list extended < identifier >
[ [ seq-# ] remark < remark-str >]
< permit | deny > < ipv4-protocol-type > < SA > < src-acl-mask > < DA > <dest-acl-mask >

      < permit | deny > tcp
         < SA > [< operator > < value >]
         < DA > [< operator > < value >]
         [established]
         [ack | fin | rst | syn]

      < permit | deny > udp
         < SA > < src-acl-mask > [< operator > < port-id >]
         < DA > < dest-acl-mask > [< operator > < port-id >]

      < permit | deny > icmp
         < SA > < src-acl-mask > < DA > < dest-acl-mask >

      < permit | deny > igmp
         < SA > < SA-mask > < DA > < dest-acl-mask > [ igmp-type ]
         [log]  (Allowed with either "deny" or "permit" ACEs.)

      . . .
      < Implicit Deny >
      exit
```

Figure 8-6. Example of General Structure Options for an Extended ACL

For example, figure 8-7 shows how to interpret the entries in an extended ACL.

```
HP Switch(config)# show running

                              ACL List Heading with
                              List Type and ID String
Running configuration:        (Name or Number)

; J8697A Configuration Editor; Created on release #YA.15.XX
      Protocol Types
hostname "HP_switch"
ip access-List extended "Sample-List-1"
   10 permit ip 10.38.130.55 0.0.0.0 10.38.130.240 0.0.0.0
   20 permit tcp 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255 eq 23
   30 remark "ALLOWS HTTP FROM SINGLE HOST."
   30 permit tcp 10.38.131.14 0.0.0.0 eq 80 0.0.0.0 255.255.255.255 eq 3871
   40 remark "DENIES HTTP FROM ANY TO ANY."    Remark for ACE at Line 40
   40 deny tcp 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255 eq 80 log
   50 deny udp 10.42.120.19 0.0.0.0 eq 69 10.38.140.44 0.0.0.0 range 3680 3690 log
   60 deny udp 0.0.0.0 255.255.255.255 10.38.99.121 0.0.0.0 log
   exit
```

Indicates all possible destination IPv4 addresses.

Denies TCP Port 80 IPv4 traffic to any destination from any source.

End-of-List Marker

ACE Action (permit or deny)

Source Addresses and Masks.
Upper entry denies certain UDP packets from a single host. Lower entry denies all UDP packets from all hosts.

Optional Source UDP or TCP Operator and Port Number
In this case, the ACL specifies UDP port 69 packets coming from the source address.

Destination Address and Mask

Optional Destination UDP or TCP Operator and Range of Port Numbers
In this case, the ACL specifies UDP port numbers 3680-3690.

Figure 8-7. Example of a Displayed Extended ACL Configuration

ACL Configuration Factors

The Sequence of Entries in an ACL Is Significant

When the switch uses an ACL to determine whether to permit or deny a packet, it compares the packet to the criteria specified in the individual Access Control Entries (ACEs) in the ACL, beginning with the first ACE in the list and proceeding sequentially until a match is found. When a match is found, the switch applies the indicated action (permit or deny) to the packet. This is significant because, once a match is found for a packet, subsequent ACEs in the same ACL will not be applied to that packet, regardless of whether they match the packet.

For example, suppose that you have applied the ACL shown in figure 8-8 to inbound IPv4 traffic on VLAN 1 (the default VLAN):

```
          Source Address      Mask    DestinationAddress    Mask

ip access-list extended "Sample-List-2"

 10 deny ip 10.28.235.10 0.0.0.0 0.0.0.0 255.255.255.255

 20 deny ip 10.28.245.89 0.0.0.0 0.0.0.0 255.255.255.255

 30 permit tcp 10.28.18.100 0.0.0.0 10.28.237.1 0.0.0.0

 40 deny tcp 10.28.18.100 0.0.0.0 0.0.0.0 255.255.255.255

 50 permit ip 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255

(Implicit Deny)  ◄─  After the last explicit ACE there is always an Implicit Deny. However, in this case it will not be used
 exit                because the last, permit ip ACL allows all IPv4 packets that earlier ACEs have not already permitted or
```

Figure 8-8. Example of a Standard ACL that Permits All IPv4 Traffic Not Implicitly Denied

Table 8-5.   Effect of the Above ACL on Inbound IPv4 Traffic in the Assigned VLAN

Line #	Action
n/a	Shows type (**extended**) and ID (Sample-List-2).
10	A packet from SA 10.28.235.10 will be denied (dropped). This ACE filters out all packets received from 10.28.235.10. As a result, IPv4 traffic from that device will not be allowed and packets from that device will not be compared against any later entries in the list.
20	A packet from SA 10.28.245.89 will be denied (dropped). This ACE filters out all packets received from 10.28.245.89. As the result, IPv4 traffic from that device will not be allowed and packets from that device will not be compared against any later entries in the list.
30	A TCP packet from SA 10.28.18.100 with a DA of 10.28.237.1 will be permitted (forwarded). Since no earlier ACEs in the list have filtered TCP packets from 10.28.18.100 and destined for 10.28.237.1, the switch will use this ACE to evaluate such packets. Any packets that meet this criteria will be forwarded. (Any packets that do not meet this TCP source-destination criteria are not affected by this ACE.)
40	A TCP packet from source address 10.28.18.100 to **any** destination address will be denied (dropped). Since, in this example, the intent is to block TCP traffic from 10.28.18.100 to any destination **except** the destination stated in the ACE at line 30, this ACE must follow the ACE at line 30. (If their relative positions were exchanged, all TCP traffic from 10.28.18.100 would be dropped, including the traffic for the 10.28.18.1 destination.)
50	Any packet from any IPv4 SA to any IPv4 DA will be permitted (forwarded). The only traffic to reach this ACE will be IPv4 packets not specifically permitted or denied by the earlier ACEs.
n/a	The *Implicit Deny* is a function the switch automatically adds as the last action in all ACLs. It denies (drops) any IPv4 traffic from any source to any destination that has not found a match with earlier entries in the ACL. In this example, the ACE at line 50 permits (forwards) any IPv4 traffic not already permitted or denied by the earlier entries in the list, so there is no traffic remaining for action by the Implicit Deny function.
exit	Marks the end of the ACL.

Allowing for the Implied Deny Function

In any ACL having one or more ACEs, there will always be a packet match. This is because the switch automatically applies an Implicit Deny as the last ACE in any ACL. This function is not visible in ACL listings, but is always present. (Refer to figure 8-8.) This means that if you configure the switch to use an ACL for filtering either inbound or outbound IPv4 traffic on a VLAN, any packets not specifically permitted or denied by the explicit entries you create will be denied by the Implicit Deny action. If you want to preempt the Implicit Deny (so that IPv4 traffic not specifically addressed by earlier ACEs in a given ACL will be permitted), insert an explicit **permit any** (for standard ACLs) or **permit ip any any** (for extended ACLs) as the last explicit ACE in the ACL.

A Configured ACL Has No Effect Until You Apply It to an Interface

The switch stores ACLs in the configuration file. Thus, until you actually assign an ACL to an interface, it is present in the configuration, but not used (and does not use any of the monitored resources described in the appendix titled "Monitored Resources" in the *Management and Configuration Guide* for your switch.)

You Can Assign an ACL Name or Number to an Interface Even if the ACL Does Not Exist in the Switch's Configuration

In this case, if you subsequently create an ACL with that name or number, the switch automatically applies each ACE as soon as you enter it in the running-config file. Similarly, if you modify an existing ACE in an ACL you already applied to an interface, the switch automatically implements the new ACE as soon as you enter it. (See "" on page 8-99.) The switch allows up to 2048 ACLs each for IPv4 and determines the total from the number of unique ACL names in the configuration. For example, if you configure two ACLs, but assign only one of them to a VLAN, the ACL total is two, for the two unique ACL names in the configuration. If you then assign the name of a nonexistent ACL to a VLAN, the new ACL total is three, because the switch now has three unique ACL names in its configuration.

(For information on switch resource use, refer to "Monitoring Shared Resources" on page 8-100. For a summary of ACL resource limits, refer to the appendix covering scalability in the latest *Management and Configuration Guide* for your switch.)

Using the CLI To Create an ACL

Command	Page
access-list (standard ACLs)	8-39
access-list (extended ACLs)	8-48

You can use either the switch CLI or an offline text editor to create an ACL. This section describes the CLI method, which is recommended for creating short ACLs. (To use the offline method, refer to "Creating or Editing ACLs Offline" on page 8-86.)

General ACE Rules

These rules apply to all IPv4 ACEs you create or edit using the CLI:

■ Inserting or adding an ACE to an ACL:

• **Named IPv4 ACLs:** Add an ACE to the end of a named ACE by using the **ip access-list** command to enter the Named ACL (**nacl**) context and entering the ACE without the sequence number. For example, if you wanted to add a "permit" ACL at the end of a list named "List-1" to allow traffic from the device at 10.10.10.100:

```
HP Switch(config)# ip access-list standard List-1
HP Switch(config-std-nacl)# permit host
10.10.10.100
```

Insert an ACE anywhere in a named ACL by specifying a sequence number. For example, if you wanted to insert a new ACE as line 15 between lines 10 and 20 in an existing ACL named "List-2" to deny IPv4 traffic from the device at 10.10.10.77:

```
HP Switch(config)# ip access-list standard List-2
HP Switch(config-std-nacl)# 15 deny host
10.10.10.77
```

• **Numbered IPv4 ACLs:** Add an ACE to the end of a numbered ACL by using the **access-list < 1 - 99 | 100 - 199 >** command. For example, if you wanted to add a "permit" ACE at the end of a list identified with the number "11" to allow IPv4 traffic from the device at 10.10.10.100:

```
HP Switch(config)# access-list 11 permit host
10.10.10.100
```

To insert an ACE *anywhere* in a numbered ACL, use the same process as described above for inserting an ACE anywhere in a *named* ACL. For example, to insert an ACE denying IPv4 traffic from the host at 10.10.10.77 as line 52 in an existing ACL identified (named) with the number 11:

```
HP Switch(config)# ip access-list standard 99
HP Switch(config-std-nacl)# 52 deny host
10.10.10.77
```

Note

After a numbered ACL has been created (using **access-list < 1 - 99 | 100 - 199 >**), it can be managed as either a named or numbered ACL, as shown above.

■ Deleting an ACE: Enter the ACL context and delete the sequence number for the unwanted ACE. (To view the sequence numbers of the ACEs in a list, use **show access-list <** *acl-name-str* **> config**.)

■ Duplicate ACEs are not allowed in the same ACL. Attempting to enter a duplicate ACE displays the **Duplicate access control entry** message.

Using CIDR Notation To Enter the IPv4 ACL Mask

You can use CIDR (Classless Inter-Domain Routing) notation to enter ACL masks. The switch interprets the bits specified with CIDR notation as the address bits in an ACL and the corresponding address bits in a packet that must match. The switch then converts the mask to inverse notation for ACL use.

Table 8-6. Examples of CIDR Notation for Masks

Address Used In an ACL with CIDR Notation	Resulting ACL Mask	Meaning
10.38.240.125/15	0.1.255.255	The leftmost 15 bits must match; the remaining bits are wildcards.
10.38.240.125/20	0.0.15.255	The leftmost 20 bits must match; the remaining bits are wildcards.
10.38.240.125/21	0.0.7.255	The leftmost 21 bits must match; the remaining bits are wildcards.
10.38.240.125/24	0.0.0.255	The leftmost 24 bits must match; the remaining bits are wildcards.
18.38.240.125/32	0.0.0.0	All bits must match.

Configuring Standard ACLs

Command Summary for Standard ACLs

Action	Command(s)	Page
Create a Standard, **Named** ACL *or* Add an ACE to the End of an Existing Standard, **Named** ACL	HP Switch(config)# ip access-list standard < *name-str* > HP Switch(config-std-nacl)# < deny | permit > < any | host <*SA*> | *SA*/< mask-length > | *SA* < *mask* >>[1] [log] [2]	8-41
Create a Standard, **Numbered** ACL or Add an ACE to the End of an Existing Standard, **Numbered** ACL	HP Switch(config)# access-list < 1-99 > < deny | permit > < any | host <*SA*> | *SA*/< mask-length > | *SA* < *mask* >> [log][2]	8-44
Use a Sequence Number To Insert an ACE in an ACL	HP Switch(config)# ip access-list standard < *name-str* | 1-99 > HP Switch(config-std-nacl)# 1-2147483647 < deny | permit > < any | host <*SA*> | *SA*/< mask-length > | *SA* < *mask* >>[1] [log] [2]	8-66
Use an ACE's Sequence Number To Delete the ACE from an ACL	HP Switch(config)# ip access-list standard < *name-str* | 1-99 > HP Switch(config-std-nacl)# no < 1-2147483647 >	8-69
Resequence the ACEs in an ACL	HP Switch(config)# ip access-list resequence < *name-str* | 1-99 > < 1-2147483647 > < 1-2147483646 >	8-70

Enter or Remove a Remark from an ACL	HP Switch(config)# ip access-list standard < *name-str* | 1-99 >
	HP Switch(config-ext-nacl)# [remark < *remark-str* > | no < 1-2147483647 > remark]
	
	For numbered, standard ACLs only, the following **remark** *commands can be substituted for the above:*
	    HP Switch(config)# access-list < 1 - 99 > remark < remark-str >
	    HP Switch(config)# [no] access-list < 1 - 99 > remark
		8-71
		8-74

Delete an ACL	HP Switch(config)# no ip access-list standard < *name-str* | 1-99 > *For numbered, standard ACLs, the following command can be substituted for the above:* HP Switch(config)# access-list < 1 - 99 > remark < remark-str >	8-64

[1]The mask can be in either dotted-decimal notation (such as 0.0.15.255) or CIDR notation (such as /20).
[2]The [log] function applies only to "deny" ACLs, and generates a message only when there is a "deny" match.

A standard ACL uses only source IPv4 addresses in its ACEs. This type of ACE is useful when you need to:

■   Permit or deny any IPv4 traffic based on source address only.

■   Quickly control the IPv4 traffic from a specific address. This allows you to isolate IPv4 traffic problems generated by a specific device, group of devices, or a subnet threatening to degrade network performance. This gives you an opportunity to troubleshoot without sacrificing performance for users outside of the problem area.

A *named*, standard ACL is identified by an alphanumeric string of up to 64 characters and is created by entering the Named ACL (**nacl**) context. A *numbered*, standard ACL is identified by a number in the range of 1 - 99 and is created without having to leave the global config context. Note that the CLI command syntax for creating a named ACL differs from the command syntax for creating a numbered ACL. For example, the first pair of entries below illustrate how to create (or enter) a named, standard ACL and enter an ACE. The next entry illustrates creating a numbered, standard ACL with the same ACE.

```
HP Switch(config)# ip access-list standard Test-List
HP Switch(config-std-nacl)# permit host 10.10.10.147

HP Switch(config)# access-list 1 permit host 10.10.10.147
```

Note that once a numbered ACL has been created, it can be accessed using the named ACL method. This is useful if it becomes necessary to edit a numbered ACL by inserting or removing individual ACEs. (Inserting or deleting an ACE is done by sequence number, and requires the Named ACL (**nacl**) context.) The switch allows a maximum of 2048 unique ACL identities (IPv4 and IPv6 combined). (For more on this topic, refer to "Monitoring Shared Resources" on page 8-100.)

Note   For a summary of standard ACL commands, refer to "Configuring Standard ACLs" on page 8-39. For a summary of all IPv4 ACL commands, refer to the command summaries on pages 8-4 and 8-5.

Configuring Named, Standard ACLs

This section describes the commands for performing the following:

■ creating and/or entering the context of a named, standard ACL

■ appending an ACE to the end of an existing list or entering the first ACE in a new list

For other IPv4 ACL topics, refer to the following:

Topic	Page
configuring numbered, standard ACLs	8-44
configuring named, extended ACLs	8-50
configuring numbered, extended ACLs	8-57
applying or removing an ACL on an interface	8-62
deleting an ACL	8-64
editing an ACL	8-65
sequence numbering in ACLs	8-66
including remarks in an ACL	8-71
displaying ACL configuration data	8-76
creating or editing ACLs offline	8-86
enabling ACL "Deny" or "Permit" logging	8-89

Entering the IPv4 "Named ACL" (nacl) Context. This command is a prerequisite to entering or editing ACEs in a named ACL.

Syntax: ip access-list standard < *name-str* >

*Places the CLI in the "Named ACL" (**nacl**) context specified by the < name-str > alphanumeric identifier. This enables entry of individual ACEs in the specified ACL. If the ACL does not already exist, this command creates it.*
*< **name-str** >: Specifies an identifier for the ACL. Consists of an alphanumeric string of up to 64 case-sensitive characters. Including spaces in the string requires that you enclose the string in single or double quotes. For example: "**Accounting ACL**".*

Refer also to table on page 8-39.

Configuring ACEs in an Named, Standard ACL. Configuring ACEs is done after using the **ip access-list standard < *name-str* >** command described above to enter the "Named ACL" (**nacl**) context of an access list. *For a standard ACL syntax summary, refer to table on page 8-39.*

Syntax: < deny | permit >

< any | host < *SA* > | *SA* <mask | *SA*/ *mask-length* >> [log]

Executing this command appends the ACE to the end of the list of ACEs in the current ACL. In the default ACL configuration, ACEs are automatically assigned consecutive sequence numbers in increments of 10 and can be renumbered using **resequence** *(page 8-70).*

Note: To insert a new ACE between two existing ACEs, precede **deny** *or* **permit** *with an appropriate sequence number. (Refer to "Inserting an ACE in an Existing ACL" on page 8-67.)*

< deny | permit >

*For named ACLs, used in the "Named ACL" (***nacl***) context to configure an ACE. Specifies whether the ACE denies or permits a packet matching the criteria in the ACE, as described below.*

< any | host < *SA* > | *SA* < *mask* > | *SA*/mask-length >

Defines the source IPv4 address (SA) a packet must carry for a match with the ACE.

- **any** — *Allows IPv4 packets from any SA.*
- **host** **< *SA* >** — *Specifies only packets having < SA > as the source. Use this criterion when you want to match the IPv4 packets from a single source address.*
- ***SA* < *mask* >** *or* ***SA* /*mask-length*** — *Specifies packets received from either a subnet or a group of IPv4 addresses. The mask format can be in either dotted-decimal format or CIDR format (number of significant bits). (Refer to "Using CIDR Notation To Enter the IPv4 ACL Mask" on page 8-38).*

  Mask Application: The mask is applied to the IPv4 address in the ACE to define which bits in a packet's SA must exactly match the SA configured in the ACE and which bits need not match. For example: **10.10.10.1/24** *and* **10.10.10.1 0.0.0.255** *both define any address in the range of 10.10.10.(1 - 255).*

  Note: Specifying a group of contiguous addresses may require more than one ACE. For more on how masks operate, refer to "How an ACE Uses a Mask To Screen Packets for Matches" on page 8-24.

[log]

This option generates an ACL log message if:
- *The action is deny or permit.*
- *There is a match.*
- *ACL logging is enabled on the switch. (Refer to "Enable ACL "Deny" or "Permit" Logging" on page 8-89.)*

(Use the debug command to direct ACL logging output to the current console session and/or to a server. Note that you must also use the **logging <** *ip-addr* **>** *command to specify the addresses of servers to which you want log messages sent. See also "Enable ACL "Deny" or "Permit" Logging" on page 8-89.)*

Example of Creating and Listing a Standard, Named ACL. This example illustrates how to create a standard, named ACL with several ACEs. This example creates an ACL that:

1. permits IPv4 traffic from a host with the address of 10.10.10.104

2. creates another ACE that blocks all other IPv4 traffic from the same subnet

3. allows all other IPv4 traffic

```
HP Switch(config)# ip access-list standard Sample-List
HP Switch(config-std-nacl)# permit host 10.10.10.104
HP Switch(config-std-nacl)# deny 10.10.10.1/24 log
HP Switch(config-std-nacl)# permit any
HP Switch(config-std-nacl)# exit
HP Switch(config)#
```

Creates the "Sample-List" ACL and enters the "Named ACL" context for this list.

Appends three ACEs to the list in the order shown.

Exits from the nacl context.

Figure 8-9. Example of Commands Used To Create an Standard, Named ACL

```
HP Switch(config)# show access-list Sample-List

Access Control Lists

  Name: Sample-List
  Type: Standard
  Applied: No

 SEQ   Entry
--------------------------------------------------------------------------------
 10    Action: permit
       IP    : 10.10.10.104      Mask: 0.0.0.0                    Note that each ACE is
                                                                  automatically assigned a
 20    Action: deny (log)                                         sequence number.
       IP    : 10.10.10.1        Mask: 0.0.0.255


 30    Action: permit
       IP    : 0.0.0.0           Mask: 255.255.255.255
```

Figure 8-10. Screen Output Listing the "Sample-List" ACL Content

Creating Numbered, Standard ACLs

Use the following general steps to create or add to an numbered, standard ACL:

1. Create a numbered, standard ACL by entering the first ACE in the list

2. Append a new ACE to the end of an existing, standard ACL

This section describes the commands for performing these steps. For other IPv4 ACL topics, refer to the following:

Topic	Page
configuring named, standard ACLs	8-41
configuring named, extended ACLs	8-50
configuring numbered, extended ACLs	8-57
applying or removing an ACL on an interface	8-62
deleting an ACL	8-64
editing an ACL	8-65
sequence numbering in ACLs	8-66
including remarks in an ACL	8-71
displaying ACL configuration data	8-76
creating or editing ACLs offline	8-86
enabling ACL "Deny" or "Permit" logging	8-89

Creating or Adding to an Standard, Numbered ACL. *This command is an alternative to using* **ip access-list standard** *< name-str > and does not use the "Named ACL" (***nacl***) context. For a standard ACL syntax summary, refer to table on page 8-39.*

Syntax: access-list < 1-99 > < deny | permit >
< any | host < *SA* > | *SA* < *mask* | *SA*/*mask-length* >> [log]

> *Appends an ACE to the end of the list of ACEs in the current IPv4 standard, numbered ACL. If the ACL does not already exist, creates both the ACL and its first ACE. In the default configuration, ACEs are automatically assigned consecutive sequence numbers in increments of 10 and can be renumbered using* **resequence** *(page 8-70).*

> *Note: To insert a new ACE between two existing ACEs in a standard, numbered ACL:*

>> *a. Use* **ip access list extended < 1 - 99 >** *to open the ACL as a named ACL.*

>> *b. Enter the desired sequence number along with the ACE keywords and variables you want.*

> *(After a numbered ACL has been created, it can be managed as either a named or numbered ACL. Refer to the "Numbered ACLs" list item on page 8-37.)*

< 1-99 >

> *Specifies the ACL identifier as a number. The switch interprets an ACL with a value in this range as a standard ACL (which filters all IPv4 traffic on the basis of SA). (To create a standard access list with an alphanumeric name (***name-str***) instead of a number, refer to "Configuring Named, Standard ACLs" on page 8-41.)*

< deny | permit >

> *Specifies whether the ACE denies or permits a packet matching the criteria in the ACE, as described next.*

< any | host < *SA* > | SA < *mask* | *SA*/*mask-length* >>

Defines the source IPv4 address (SA) a packet must carry for a match with the ACE.

- **any** — *Allows IPv4 packets from any SA.*
- **host < *SA* >** — *Specifies only packets having < *SA* > as the source. Use this criterion when you want to match only the IPv4 packets from a single SA.*
- ***SA* < *mask* >** *or **SA /*mask-length*** — Specifies packets received from an SA, where the SA is either a subnet or a group of IPv4 addresses. The mask format can be in either dotted-decimal format or CIDR format (number of significant bits). (Refer to "Using CIDR Notation To Enter the IPv4 ACL Mask" on page 8-38).*

  SA Mask Application: *The mask is applied to the SA in the ACE to define which bits in a packet's SA must exactly match the SA configured in the ACL and which bits need not match.*

  Example: **10.10.10.1/24** *and* **10.10.10.1 0.0.0.255** *both define any address in the range of 10.10.10.(1 - 255).*

  Note: *Specifying a group of contiguous addresses may require more than one ACE. For more on how masks operate in ACLs, refer to "How an ACE Uses a Mask To Screen Packets for Matches" on page 8-24.*

[log]

This option can be used after the SA to generate an Event Log message if:

- *The action is **deny** or **permit**.*
- *There is a match.*
- *ACL logging is enabled. (Refer to "Enabling ACL Logging on the Switch" on page 8-91.)*

Example of Creating and Viewing a Standard ACL. This example cre-
ates a standard, numbered ACL with the same ACE content as show in figure
8-9 on page 8-43.

```
HP Switch(config)# access-list 17 permit host 10.10.10.104
HP Switch(config)# access-list 17 deny 10.10.10.1/24 log
HP Switch(config)# access-list 17 permit any
HP Switch(config)# show access-list 17

Access Control Lists

  Name: 17
  Type: Standard
  Applied: No

 SEQ  Entry
-------------------------------------------------------------------------------
 10   Action: permit                                        Note that each ACE is
      IP    : 10.10.10.104      Mask: 0.0.0.0               automatically assigned a
                                                            sequence number.
 20   Action: deny (log)
      IP    : 10.10.10.1        Mask: 0.0.0.255

 30   Action: permit
      IP    : 0.0.0.0           Mask: 255.255.255.255
```

Figure 8-11. Standard, Numbered ACL with the Same ACEs as the Standard, Named ACL in Figure 8-9

Configuring Extended ACLs

Command Summary for Extended ACLs

Action	Command(s)	Page
Create an Extended, **Named** ACL *or* Add an ACE to the End of an Existing, Extended ACL	HP Switch(config)# ip access-list extended < *name-str* | 100-199 > HP Switch(config-std-nacl)# < deny | permit > < ip | *ip-protocol* | *ip-protocol-nbr* > < any | host <*SA* > | *SA*/< mask-length > | *SA* < *mask* >>[1] < any | host < *DA* > | *DA*/< mask-length > | *DA* < *mask* >>[1]	8-50

	[tcp | udp]
	      < any | host <*SA* > | *SA*/< mask-length > | *SA* < *mask* >>[1]
	        [*comparison-operator* < *value* >]]
	      < any | host <*DA* > | *DA*/< mask-length > | *DA* < *mask* >>[1]
	        [*comparison-operator* < *value* >]
	[established]
	[ack] [fin] [rst] [syn][2]
		8-50

	< igmp > < any | host <*SA* > | *SA*/< mask-length > | *SA* < *mask* >>[1] < any | host < *DA* > | *DA*/< mask-length > | *DA* < *mask* >>[1] [*igmp-packet-type*]	8-50
	< icmp > < any | host <*SA* > | *SA*/< mask-length > | *SA* < *mask* >>[1] < any | host < *DA* > | *DA*/< mask-length > | *DA* < *mask* >>[1] [log][2]	8-50

Create an Extended, **Numbered** ACL
 	
or
 	
Add an ACE to the End of an Existing, **Numbered** ACL
	HP Switch(config)# access-list < 100-199 > < deny | permit >
	    < *ip-options* | *tcp/udp-options* | *igmp-options* | *icmp-options* >
	    [log][2]
	    Note: *Uses the same IPv4, TCP/UDP, IGMP, and ICMP options as shown above for "Create an Extended, Named ACL".*
		8-57

Insert an ACE by Assigning a Sequence Number	HP Switch(config)# ip access-list extended < *name-str* | 100-199 > HP Switch(config-ext-nacl)# 1-2147483647 < deny | permit > *Uses the options shown above for "Create an Extended, Named ACL".*	8-67
Use Sequence Number To Delete an ACE	HP Switch(config)# ip access-list extended < *name-str* | 100-199 > HP Switch(config-std-nacl)# no < 1-2147483647 >	8-69
Resequence the ACEs in an ACL	HP Switch(config)# ip access-list resequence < *name-str* | 100-199 > < 1-2147483647 > < 1-2147483646 >	8-70

Action	Command(s)	Page
Enter or Remove a Remark	HP Switch(config)# ip access-list extended < *name-str* | 100-199 > HP Switch(config-ext-nacl)# [remark < *remark-str* > | no < 1 - 2147483647 > remark]	8-71 8-74
	For numbered, extended ACLs only, the following **remark** *commands can be substituted for the above:* HP Switch(config)# access-list < 100 - 199 > remark < remark-str > HP Switch(config)# [no] access-list < 100 - 199 > remark	
Delete an Extended ACL	HP Switch(config)# no ip access-list extended < *name-str* | 100-199 >	8-64
	For numbered, extended ACLs only, the following command can also be used: HP Switch(config)# no access-list < 100 - 199 >	

[1]The mask can be in either dotted-decimal notation (such as 0.0.15.255) or CIDR notation (such as /20).
[2]TCP flag (control bit) options for destination TCP.

Standard ACLs use only source IPv4 addresses for filtering criteria, extended ACLs use multiple filtering criteria. This enables you to more closely define your IPv4 packet-filtering. Extended ACLs enable filtering on the following:

■ Source and destination IPv4 addresses (required), in one of the following options:
  • specific host
  • subnet or group of addresses
  • any address

■ choice of any IPv4 protocol

■ optional packet-type criteria for IGMP, and ICMP traffic

■ optional source and/or destination TCP or UDP port, with a further option for comparison operators and (for TCP) an option for establishing connections

■ filtering for TCP traffic based on either TCP control bits or whether the subject traffic is initiating a connection ("established" option)

The switch allows up to 2048 ACLs in any combination of IPv4 and IPv6 ACLs, and determines the total from the number of unique identifiers in the configuration. For example, configuring two ACLs results in an ACL total of two, even if neither is assigned to an interface. If you then assign a nonexistent ACL to an interface, the new ACL total is three, because the switch now has three unique ACL names in its configuration. (For more on ACL limits, refer to "Monitoring Shared Resources" on page 8-100.)

Configuring Named, Extended ACLs

For a match to occur with an ACE in an extended ACL, a packet must have the source and destination address criteria specified by the ACE, as well as any IPv4 protocol-specific criteria included in the command.

Use the following general steps to create or add to a named, extended ACL:

1.   Create and/or enter the context of a named, extended ACL.

2.   Enter the first ACE in a new, extended ACL or append an ACE to the end of an existing, extended ACL.

This section describes the commands for performing these steps. For other ACL topics, refer to the following:

Creating a Named, Extended ACL and/or Entering the "Named ACL" (nacl) Context. This command is a prerequisite to entering or editing ACEs in a named, extended ACL. (For a summary of the extended ACL syntax options, refer to table on page 8-48.)

Syntax: ip access-list extended < *name-str* >

> *Places the CLI in the "Named ACL" (**nacl**) context specified by the < name-str > alphanumeric identifier. This enables entry of individual ACEs in the specified ACL. If the ACL does not already exist, this command creates it.*
> *< **name-str** >: Specifies an alphanumeric identifier for the ACL. Consists of an alphanumeric string of up to 64 case-sensitive characters. Including spaces in the string requires that you enclose the string in single or double quotes. For example: "**Accounting ACL**". You can also use this command to access an existing, numbered ACL. Refer to "Using the CLI To Edit ACLs" on page 8-65*

```
HP Switch(config)# ip access-list extended Sample-List
HP Switch(config-ext-nacl)#
```

Figure 8-12. Example of Entering the Named ACL Context

Configure ACEs in a Named, Extended ACL and/or Enter the "Named ACL" (nacl) Context. Configuring ACEs is done after using the **ip access-list standard** *< name-str >* command described on page 8-51 to enter the "Named ACL" (**nacl**) context of an ACL. For an extended ACL syntax summary, refer to table on page 8-48.

Syntax: < deny | permit > < ip | *ip-protocol* | *ip-protocol-nbr* >
(nacl < any | host < *SA* > | *SA/ mask-length* | *SA* < *mask* > >
context) < any | host < *DA* > | *DA/ mask-length* | *DA* < *mask* > >
[log]

Appends an ACE to the end of the list of ACEs in the current ACL. In the default configuration, ACEs are automatically assigned consecutive sequence numbers in increments of 10 and can be renumbered using **resequence** *(page 8-70).*

Note: To insert a new ACE between two existing ACEs in an extended, named ACL, precede **deny** *or* **permit** *with an appropriate sequence number along with the ACE keywords and variables you want. (Refer to "Inserting an ACE in an Existing ACL" on page 8-67.)*

For a match to occur, a packet must have the source and destination addressing criteria specified in the ACE, as well as:

- *the protocol-specific criteria configured in the ACE, including any included, optional elements (described later in this section)*

< deny | permit >

*For named ACLs, these keywords are used in the "Named ACL" (***nacl***) context to specify whether the ACE denies or permits a packet matching the criteria in the ACE, as described below.*

< ip | *ip-protocol* | *ip-protocol-nbr* >

Used after **deny** *or* **permit** *to specify the packet protocol type required for a match. An extended ACL must include one of the following:*

- **ip** — *any IPv4 packet.*
- *ip-protocol* — *any one of the following IPv4 protocol names:*

ip-in-ip	ipv6-in-ip	gre	esp	ah
ospf	pim	vrrp	sctp	tcp*
udp*	icmp*	igmp*		

- *ip-protocol-nbr* — *the protocol number of an IPv4 packet type, such as "8" for Exterior Gateway Protocol or 121 for Simple Message Protocol. (For a listing of IPv4 protocol numbers and their corresponding protocol names, refer to the IANA "Protocol Number Assignment Services" at www.iana.com.) (Range: 0 - 255)*

< any | host < *SA* > | *SA* < *mask* > | *SA/ mask-length*

This is the first instance of IPv4 addressing in an extended ACE. It follows the protocol specifier and defines the <u>source address</u> (SA) a packet must carry for a match with the ACE.

- **any** — *Allows IPv4 packets from any SA.*
- **host** < *SA* > — *Specifies only packets having a single address as the SA. Use this criterion when you want to match only the IPv4 packets from a single SA.*
- *SA < mask >* or *SA/mask-length* — *Specifies packets received from an SA, where the SA is either a subnet or a group of addresses. The mask can be in either dotted-decimal format or CIDR format (number of significant bits). Refer to "Using CIDR Notation To Enter the IPv4 ACL Mask" on page 8-38.*

   SA Mask Application: The mask is applied to the SA in the ACL to define which bits in a packet's SA must exactly match the SA configured in the ACL and which bits need not match.

   Example: 10.10.10.1/24 and 10.10.10.1 0.0.0.255 both define any address in the range of 10.10.10.(1 - 255).

   Note: Specifying a group of contiguous addresses may require more than one ACE. For more on how masks operate in ACLs, refer to "How an ACE Uses a Mask To Screen Packets for Matches" on page 8-24.

< any | host < *DA* > | *DA/mask*-length | *DA/* < *mask* >>

This is the second instance of IPv4 addressing in an extended ACE. It follows the first (SA) instance, described earlier, and defines the destination address (DA) that a packet must carry in order to have a match with the ACE.

- **any** — *Allows routed IPv4 packets to any DA.*
- **host** < *DA* > — *Specifies only packets having* **DA** *as the destination address. Use this criterion when you want to match only the IPv4 packets for a single DA.*
- ***DA/mask-length*** *or* **DA< mask >** — *Specifies packets intended for a destination address, where the address is either a subnet or a group of addresses. The mask format can be in either dotted-decimal format or CIDR format (number of significant bits). Refer to "Using CIDR Notation To Enter the IPv4 ACL Mask" on page 8-38.*

  DA Mask Application: *The mask is applied to the DA in the ACL to define which bits in a packet's DA must exactly match the DA configured in the ACL and which bits need not match. See also the above example and note.*

[log]

This option can be used after the DA to generate an Event Log message if:

- *The action is* **deny** *or* **permit***.*
- *There is a match.*
- *ACL logging is enabled. (Refer to "Enabling ACL Logging on the Switch" on page 8-91.)*

Options for TCP and UDP Traffic in Extended ACLs. An ACE designed
to permit or deny TCP or UDP traffic can optionally include port number
criteria for either the source or destination, or both. Use of TCP criteria also
allows the **established** option for controlling TCP connection traffic. (For a
summary of the extended ACL syntax options, refer to table on page 8-48.)

Syntax: < deny | permit > tcp
< *SA* > [*comparison-operator* < *tcp-src-port* >]
< *DA* > [*comparison-operator* < *tcp-dest-port* >]
[established]
[ack] [fin] [rst] [syn]

Syntax: < deny | permit > udp
< *SA* > [*comparison-operator* < *udp-src-port* >]
< *DA* > [*comparison-operator* < *udp-dest-port* >]

*In an extended ACL using either **tcp** or **udp** as the packet protocol
type, you can optionally use TCP or UDP source and/or desti-
nation port numbers or ranges of numbers to further define the
criteria for a match. For example:*

```
#deny tcp host 10.20.10.17 eq 23 host 10.20.10.155
 established
#permit tcp host 10.10.10.100 host 10.20.10.17
  eq telnet
#deny udp 10.30.10.1/24 host 10.20.10.17 range
 161 162
```

[*comparison-operator* < *tcp/udp-src-port* >]

*To specify a TCP or UDP <u>source</u> port number in an ACE, (1)
select a comparison operator from the following list and (2)
enter the port number or a well-known port name.*

Comparison Operators:

- **eq** < ***tcp/udp-port-nbr*** > — *"Equal To"; to have a match with the ACE entry, the TCP or UDP source port number in a packet must be equal to* < ***tcp/udp-port-nbr*** >.
- **gt** < ***tcp/udp-port-nbr*** > — *"Greater Than"; to have a match with the ACE entry, the TCP or UDP source port number in a packet must be greater than* < ***tcp/udp-port-nbr*** >.
- **lt** < ***tcp/udp-port-nbr*** > — *"Less Than"; to have a match with the ACE entry, the TCP or UDP source port number in a packet must be less than* < ***tcp/udp-port-nbr*** >.
- **neq** < ***tcp/udp-port-nbr*** > — *"Not Equal"; to have a match with the ACE entry, the TCP or UDP source port number in a packet must not be equal to* **< *tcp/udp-port-nbr* >**.
- **range** **< *start-port-nbr* > < *end-port-nbr* >** — *For a match with the ACE entry, the TCP or UDP source-port number in a packet must be in the range* <**start-port-nbr** > **< *end-port-nbr* >**.

Port Number or Well-Known Port Name:

Use the TCP or UDP port number required by your application. The switch also accepts these well-known TCP or UDP port names as an alternative to their port numbers:

- ***TCP***: bgp, dns, ftp, http, imap4, ldap, nntp, pop2, pop3, smtp, ssl, telnet
- ***UDP***: bootpc, bootps, dns, ntp, radius, radius-old, rip, snmp, snmp-trap, tftp

To list the above names, press the **[Shift] [?]** *key combination after entering an operator. For a comprehensive listing of port numbers, visit www.iana.org/assignments/port-numbers.*

[*comparison-operator* < *tcp-dest-port* >] [established]
[*comparison-operator* < *udp-dest-port* >]

This option, if used, is entered immediately after the **< DA >** *entry. To specify a TCP or UDP port number, (1) select a comparison operator and (2) enter the port number or a well-known port name.*

Comparison Operators and Well-Known Port Names —
These are the same as are used with the TCP/UDP source-port options, and are listed earlier in this command description.

[established] — *This option applies only where TCP is the configured protocol type. It blocks the synchronizing packet associated with establishing a TCP connection in one direction on a VLAN while allowing all other IPv4 traffic for the same type of connection in the opposite direction. For example, a Telnet connect requires TCP traffic to move both ways between a host and the target device. Simply applying a* **deny** *to inbound Telnet traffic on a VLAN would prevent Telnet sessions in either direction because responses to outbound requests would be blocked. However, by using the* **established** *option, inbound Telnet traffic arriving in response to outbound Telnet requests would be permitted, but inbound Telnet traffic trying to establish a connection would be denied.*

TCP Control Bits. *In a given ACE for filtering TCP traffic you can configure one or more of these options:*

[ack] — *Acknowledgement.*

[fin] — *Sender finished.*

[rst] — *Connection reset.*

[syn] — *TCP control bit: sequence number synchronize.*

For more on using TCP control bits, refer to RFC 793.

Configuring Numbered, Extended ACLs

This section describes the commands for performing the following in a numbered, extended ACL:

- creating the ACL by entering the first ACE in the list

- appending a new ACE to the end of an existing ACL

For other IPv4 ACL topics, refer to the following:

Topic	Page
configuring named, standard ACLs	8-41
configuring numbered, standard ACLs	8-44
configuring named, extended ACLs	8-50
applying or removing an ACL on an interface	8-62
deleting an ACL	8-64

Creating or Adding to an Extended, Numbered ACL. This command is an alternative to using **ip access-list extended < *name-str* >** and does not use the Named ACL (**nacl**) context. (For an extended ACL syntax summary, refer to table on page 8-48.)

Syntax: access-list < 100-199 > < deny | permit > < ip | *ip-protocol* | *ip-protocol-nbr* >
< any | host < *SA* > | *SA*/mask-length | SA < *mask* >>
< any | host < *DA* > | *DA*/mask-length | DA < *mask* >>
[log]

> *If the ACL does not already exist, this command creates the specified ACL and its first ACE. If the ACL already exists, the new ACE is appended to the end of the configured list of explicit ACEs. In the default configuration, the ACEs in an ACL will automatically be assigned consecutive sequence numbers in increments of 10 and can be renumbered with* **resequence** *(page 8-70).*
>
> *Note: To insert a new ACE between two existing ACEs in an extended, numbered ACL:*
>
> > *a.   Use* **ip access list extended < 100 - 199 >** *to open the ACL as a named ACL.*
> >
> > *b.   Enter the desired sequence number along with the ACE statement you want.*
>
> *(Refer to the "Numbered ACLs" list item on page 8-37.)*
> *For a match to occur, a packet must have the source and destination addressing criteria specified in the ACE.*

> < 100-199 >
>
> *Specifies the ACL ID number. The switch interprets a numeric ACL with a value in this range as an extended ACL.*

< deny | permit >

*Specifies whether to deny (**drop**) or permit (forward) a packet that matches the criteria specified in the ACE, as described below.*

< ip | *ip-protocol* | *ip-protocol-nbr* >

Specifies the packet protocol type required for a match. An extended ACL must include one of the following:

- **ip** — *any IPv4 packet.*
- **ip-protocol** — *any one of the following IPv4 protocol names:*

ip-in-ip	**ipv6-in-ip**	**gre**	**esp**	**ah**
ospf	**pim**	**vrrp**	**sctp**	**tcp***
udp*	**icmp***	**igmp***		

- **ip-protocol-nbr** — *the protocol number of an IPv4 packet type, such as "8" for Exterior Gateway Protocol or 121 for Simple Message Protocol. (For a listing of IPv4 protocol numbers and their corresponding protocol names, refer to the IANA "Protocol Number Assignment Services" at www.iana.com.) (Range: 0 - 255)*

* For TCP, UDP, ICMP, and IGMP, additional criteria can be specified, as described later in this section.

< any | host < *SA* > | *SA/mask-length I SA* < *mask* >>

In an extended ACL, this parameter defines the source address (SA) that a packet must carry in order to have a match with the ACE.

- **any** — *Specifies all inbound IPv4 packets.*
- **host** < *SA* > — *Specifies only inbound IPv4 packets from a single address. Use this option when you want to match only the IPv4 packets from a single source address.*
- **SA/mask-length** *or* **SA < mask >** — *Specifies packets received from an SA, where the SA is either a subnet or a group of IPv4 addresses. The mask can be in either dotted-decimal format or CIDR format with the number of significant bits. Refer to "Using CIDR Notation To Enter the IPv4 ACL Mask" on page 8-38.*

SA Mask Application: The mask is applied to the SA in the ACL to define which bits in a packet's source SA must exactly match the address configured in the ACL and which bits need not match.

> *Example: 10.10.10.1/24 and 10.10.10.1 0.0.0.255 both define any IPv4 address in the range of 10.10.10.(1-255).*

Note: Specifying a group of contiguous IPv4 addresses may require more than one ACE. For more on how masks operate in ACLs, refer to "How an ACE Uses a Mask To Screen Packets for Matches" on page 8-24.

< any | host < *DA* > | *DA/mask*-length >

This is the second instance of addressing in an extended ACE. It follows the first (SA) instance, described earlier, and defines the destination address (DA) that a packet must carry in order to have a match with the ACE. The options are the same as shown for < SA >.

- **any** — *Allows routed IPv4 packets to any DA.*
- **host** < *DA* > — *Specifies only the packets having **DA** as the destination address. Use this criterion when you want to match only the IPv4 packets for a single DA.*
- ***DA/mask-length** or **DA < mask > — Specifies packets intended for a destination address, where the address is either a subnet or a group of IPv4 addresses. The mask format can be in either dotted-decimal format or CIDR format (number of significant bits). Refer to "Using CIDR Notation To Enter the IPv4 ACL Mask" on page 8-38.*

DA Mask Application: The mask is applied to the DA in the ACL to define which bits in a packet's DA must exactly match the DA configured in the ACL and which bits need not match. See also the above example and note.

[log]

Optional; generates an Event Log message if:

- *The action is **deny** or **permit**.*
- *There is a match.*
- *ACL logging is enabled on the switch. (Refer to "Enabling ACL Logging on the Switch" on page 8-91)*

Additional Options for TCP and UDP Traffic. An ACE designed to permit or deny TCP or UDP traffic can optionally include port number criteria for either the source or destination, or both. Use of TCP criteria also allows the **established** option for controlling TCP connection traffic. (For a summary of the extended ACL syntax options, refer to "Command Summary for Extended ACLs" on page 8-48 or "Options for TCP and UDP Traffic in Extended ACLs" on page 8-55.)

Syntax: access-list < 100 - 199 > < deny | permit > < tcp | udp >

< *SA* > [*comparison-operator* < *tcp/udp-src-port* >]

< *DA* > [*comparison-operator* < *tcp-dest-port* >] [established]
< *DA* > [*comparison-operator* < *udp-dest-port* >]

This source-port and destination-port TCP/UDP criteria is identical to the criteria described for TCP/UDP use in named, extended ACLs, beginning on page 8-55.

Adding or Removing an ACL Assignment On an Interface

Filtering IPv4 Traffic Inbound on a VLAN

For a given VLAN interface, you can assign an ACL as a VACL to filter any IPv4 traffic entering the switch on that VLAN. You can also use the same ACL for assignment to multiple VLANs. For limits and operating rules, refer to "IPv4 ACL Configuration and Operating Rules" on page 8-22.

Syntax: [no] vlan < *vid* > ip access-group < *identifier* > vlan
where: < *identifier* > = either a ACL name or an ACL ID number.

> *Assigns an ACL as a VACL to a VLAN to filter any IPv4 traffic entering the switch on that VLAN. You can use either the global configuration level or the VLAN context level to assign or remove a VACL.*
>
> ***Note:*** *The switch allows you to assign a nonexistent ACL name or number to a VLAN. In this case, if you subsequently configure an ACL with that name or number, it automatically becomes active on the assigned VLAN. Also, if you delete an assigned ACL from the switch without subsequently using the "***no***" form of this command to remove the assignment to a VLAN, the ACL assignment remains and will automatically activate any new ACL you create with the same identifier (name or number).*

```
HP Switch(config)# vlan 20 ip access-group My-List vlan      Enables a VACL from the
                                                             Global Configuration
                                                             Level
HP Switch(config)# vlan 20
HP Switch(vlan-20)# ip access-group 155 vlan
HP Switch(vlan-20)# exit                                     Enables a VACL from a
                                                             VLAN Context.

HP Switch(config)# no vlan 20 ip access-group My-List vlan   Disables a VACL from the
                                                             Global Configuration
                                                             Level
HP Switch(config)# vlan 20
HP Switch(vlan-20)# no ip access-group 155 vlan
HP Switch(vlan-20)# exit                                     Disables a VACL from a
                                                             VLAN Context.
```
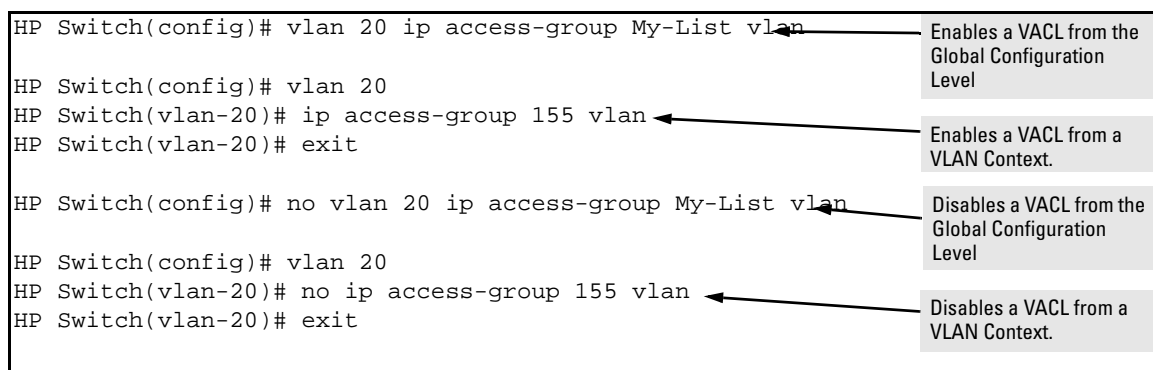
Figure 8-13. Methods for Enabling and Disabling VACLs

Filtering Inbound IPv4 Traffic Per Port

For a given port, port list, or static port trunk, you can assign an ACL as a static port ACL to filter any IPv4 traffic entering the switch on that interface. You can also use the same ACL for assignment to multiple interfaces. For limits and operating rules, refer to "IPv4 ACL Configuration and Operating Rules" on page 8-22.

Syntax: [no] interface < *port-list* | Trk*x* > ip access-group < *identifier* > in
where: < *identifier* > = either a ACL name or an ACL ID number.

> *Assigns an ACL as a static port ACL to a port, port list, or static trunk to filter any IPv4 traffic entering the switch on that interface. You can use either the global configuration level or the interface context level to assign or remove a static port ACL.*
>
> ***Note:*** *The switch allows you to assign a nonexistent ACL name or number to an interface. In this case, if you subsequently configure an ACL with that name or number, it automatically becomes active on the assigned interface. Also, if you delete an assigned ACL from the switch without subsequently using the "**no**" form of this command to remove the assignment to an interface, the ACL assignment remains and will automatically activate any new ACL you create with the same identifier (name or number).*

```
HP Switch(config)# interface 10 ip access-group My-List in       Enables a static port ACL
                                                                 from the Global
                                                                 Configuration level.
HP Switch(config)# interface 10
HP Switch(eth-10)# ip access-group 155 in
HP Switch(eth-10)# exit                                          Enables a static port ACL
                                                                 from a port context.

HP Switch(config)# no interface 10 ip access-group My-List in    Disables a static port ACL
                                                                 from the Global
                                                                 Configuration level.
HP Switch(config)# interface 10
HP Switch(eth-10)# no ip access-group 155 in                     Uses a VLAN context to
HP Switch(eth-10)# exit                                          disable a static port ACL.
```

Figure 8-14. Methods for Enabling and Disabling ACLs

Deleting an ACL

Syntax: no ip access-list standard < *name-str* | 1-99 >

no ip access-list extended < *name-str* | 100-199 >

no access-list < 1 - 99 | 100 - 199 >

> *Removes the specified ACL from the switch's running-config file.*

> ***Notes:*** *If an ACL name is assigned to an interface before the ACL itself has actually been created, then the switch creates an "empty" version of the ACL in the running configuration and assigns the empty ACL to the interface. Subsequently populating the empty ACL with explicit ACEs causes the switch to automatically activate the ACEs as they are created and to implement the implicit deny at the end of the ACL.*

> *Deleting an ACL from the running configuration while the ACL is currently assigned on an interface results in an "empty" version of the ACL in the running configuration and on the interface. Subsequently removing the ACL from the interface also removes the empty ACL from the running configuration.*

> *If you need to remove an ACL identifier assignment on an interface, refer to "Adding or Removing an ACL Assignment On an Interface" on page 8-62*

Editing an Existing ACL

The CLI provides the capability for editing in the switch by using sequence numbers to insert or delete individual ACEs. An offline method is also available. This section describes using the CLI for editing ACLs. To use the offline method for editing ACLs, refer to "Creating or Editing ACLs Offline" on page 8-86.

Using the CLI To Edit ACLs

You can use the CLI to delete individual ACEs from anywhere in an ACL, append new ACEs to the end of an ACL, and insert new ACEs anywhere within an ACL.

General Editing Rules

- **Named ACLs:**
  - When you enter a new ACE in a named ACL without specifying a sequence number, the switch inserts the ACE as the last entry in the ACL.
  - When you enter a new ACE in a named ACL and include a sequence number, the switch inserts the ACE according to the position of the sequence number in the current list of ACEs.

- **Numbered ACLs:** When using the **access-list < 1 - 99 | 100 - 199 >** command to create or add ACEs to a numbered ACL, each new ACE you enter is added to the end of the current list. (This command does not offer a **< seq-# >** option for including a sequence number to enable inserting an ACE at other points in the list.) Note, however, that once a numbered list has been created, you have the option of accessing it in the same way as a named list by using the **ip access-list < standard | extended >** command. This enables you to edit a numbered list in the same way that you would edit a named list. (See the next item in this list.)

■ You can delete any ACE from any ACL (named or numbered) by using
   the **ip access-list** command to enter the ACL's context, and then using
   the **no < *seq-#* >** command (page 8-69).

■ Deleting the last ACE from an ACL leaves the ACL in memory. In this
   case, the ACL is "empty" and cannot perform any filtering tasks. (In
   any ACL the Implicit Deny does not apply unless the ACL includes at
   least one explicit ACE.)

Sequence Numbering in ACLs

The ACEs in any ACL are sequentially numbered. In the default state, the
sequence number of the first ACE in a list is "10" and subsequent ACEs are
numbered in increments of 10. For example, the following **show run** output
lists three ACEs with default numbering in a list named "My-List":

```
ip access-list standard "My-List"
   10 permit 10.10.10.25 0.0.0.0
   20 permit 10.20.10.117 0.0.0.0
   30 deny 10.20.10.1 0.0.0.255
   exit
```

Figure 8-15. Example of the Default Sequential Numbering for ACEs

You can add an ACE to the end of a named or numbered ACL by using either
access-list for numbered ACLs or **ip access-list** for named ACLs:

```
HP Switch(config)# access-list 2 permit any        Appends an ACE to
                                                   the end of a standard,
                                                   numbered ACL.


 HP Switch(config)# ip access-list standard My-list
 HP Switch(config-ext-nacl)# permit ip any host 10.10.10.125


                                     Enters the context of an extended ACL
                                     and appends an ACE to the end of the list.
```

Figure 8-16. Examples of Adding an ACE to the end of Numbered or Named ACLs

For example, to append a fourth ACE to the end of the ACL in figure 8-15:

```
HP Switch(config)# ip access-list standard My-List
HP Switch(config-std-nacl)# permit any
HP Switch(config-std-nacl)# show run
.
.
.
ip access-list standard "My-List"
   10 permit 10.10.10.25 0.0.0.0
   20 permit 10.20.10.117 0.0.0.0
   30 deny 10.20.10.1 0.0.0.255
   40 permit 0.0.0.0 255.255.255.255
   exit
```

Figure 8-17. Example of Appending an ACE to an Existing List

Note

When using the **access-list < 1 - 99 | 100 - 199 > < permit | deny > < SA >** command to create an ACE for a numbered ACL, the ACE is always added to the end of the current list and given the appropriate sequence number. However, once a numbered list has been created, you can use the **ip access-list** command to open it as a named ACL and specify a nondefault sequence number, as described in the next section.

Inserting an ACE in an Existing ACL

This action uses a sequence number to specify where to insert a new ACE into an existing sequence of ACLs.

Syntax: ip access-list < standard | extended > < *name-str* | 1 - 99 | 100 - 199 >

<1-2147483647> < permit | deny > < *standard-acl-ip-criteria* > [log]
<1-2147483647> < permit | deny > < *extended-acl-ip-criteria* > [*options*]

The first command enters the "Named-ACL" context for the specified ACL. The remaining two commands insert a new ACE in a standard or extended ACL, respectively. (For details on these criteria and options, refer to table , "Command Summary for Standard IPv4 ACLs" —page 8-4, and table , "Command Summary for IPv4 Extended ACLs" —page 8-5.)

Note: Entering an ACE that would result in an out-of-range sequence number is not allowed. Use the resequence command to free up ACE numbering availability in the ACL. Refer to "Resequencing the ACEs in an ACL" on page 8-70.

To insert a new ACE between existing ACEs in a list:

1. Use **ip access-list** to enter the "Named-ACL" (**nacl**) context of the ACE. *This applies regardless of whether the ACE was originally created as a numbered ACL or a named ACL.*

2. Begin the ACE command with a sequence number that identifies the position you want the ACE to occupy. (The sequence number range is 1-2147483647.)

3. Complete the ACE with the command syntax appropriate for the type of ACL you are editing.

For example, inserting a new ACE between the ACEs numbered 10 and 20 in figure 8-17 requires a sequence number in the range of 11-19 for the new ACE.

```
HP Switch(config)# ip access-list standard My-List
HP Switch(config-std-nacl)# 15 deny 10.10.10.1/24
HP Switch(config-std-nacl)# show run
.
.
.
ip access-list standard "My-List"
   10 permit 10.10.10.25 0.0.0.0
   15 deny 10.10.10.1 0.0.0.255
   20 permit 10.20.10.117 0.0.0.0
   30 deny 10.20.10.1 0.0.0.255
   40 permit 0.0.0.0 255.255.255.255
   exit
```

Enters the "Named-ACL context for "My-List".

Inserts the new ACE.
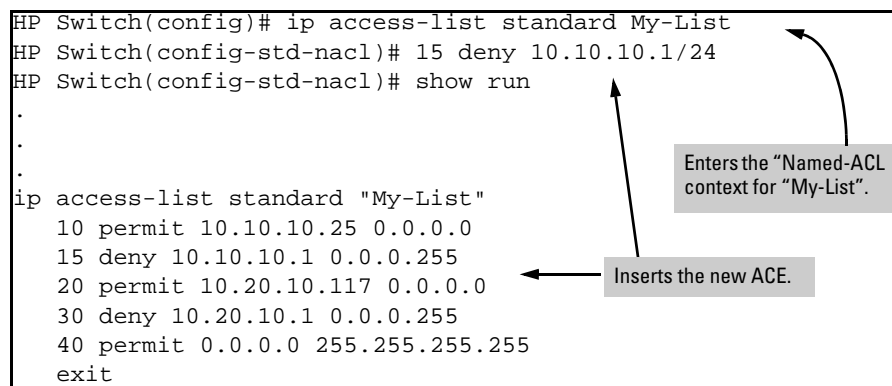
Figure 8-18. Example of Inserting an ACE in an Existing ACL

In the following example, the first two ACEs entered become lines 10 and 20 in the list. The third ACE entered is configured with a sequence number of 15 and is inserted between lines 10 and 20.

```
HP Switch(config)# ip access-list standard List-01
HP Switch(config-std-nacl)# permit 10.10.10.1/24       Becomes Line 10
HP Switch(config-std-nacl)# deny 10.10.1.1/16          Becomes Line 20
HP Switch(config-std-nacl)# 15 permit 10.10.20.1/24
HP Switch(config-std-nacl)# show run

                                                       Lines 10 and 20
Running configuration:                                 were automatically
. . .                                                  numbered accord-
ip access-list standard "List-01"                      ing to their order of
   10 permit 10.10.10.1 0.0.0.255                       entry in the list. Line
   15 permit 10.10.20.1 0.0.0.255                       15 was explicitly
   20 deny 10.10.1.1 0.0.255.255                        numbered by the 15
   exit                                                 permit command
                                                       and was inserted in
                                                       its proper place in
                                                       the list.
```

Figure 8-19. Example of Inserting an ACE into an Existing Sequence

Deleting an ACE from an Existing ACL

This action uses ACL sequence numbers to delete ACEs from an ACL.

Syntax: ip access-list < standard | extended > < *name-str* | 1 - 99 | 100 - 199 >

no < *seq-#* >

The first command enters the "Named-ACL" context for the specified ACL. The **no** *command deletes the ACE corresponding to the sequence number entered. (Range: 1 - 2147483647)*

1. To find the sequence number of the ACE you want to delete, use **show run** or **show access-list < *name-str* | 1 - 99 | 100-199 >** to view the ACL.

2. Use **ip access-list** to enter the "Named-ACL" (**nacl**) context of the ACE. *This applies regardless of whether the ACE was originally created as a numbered ACL or a named ACL.*

3. In the "Named-ACL" context, type **no** and enter the sequence number of the ACE you want to delete.

Figure 8-20 illustrates the process for deleting an ACE from a list:

```
HP Switch(config)# show run
. . .
ACL Before Deleting an ACE
ip access-list standard "My-List"
   10 permit 10.10.10.25 0.0.0.0
   15 deny 10.10.10.1 0.0.0.255
   20 permit 10.20.10.117 0.0.0.0
   30 deny 10.20.10.1 0.0.0.255                This command enters the
   40 permit 0.0.0.0 255.255.255.255           "Named-ACL" (nacl)
   exit                                        context for "My-List".
HP Switch(config)# ip access-list standard My-List
HP Switch(config-std-nacl)# no 20             This command deletes the
HP Switch(config-std-nacl)# show run          ACE at line 20.
. . .
ACL After Deleting the ACE at Line 20
ip access-list standard "My-List"
   10 permit 10.10.10.25 0.0.0.0
   15 deny 10.10.10.1 0.0.0.255                The ACE at line 20 has been
   30 deny 10.20.10.1 0.0.0.255                removed.
   40 permit 0.0.0.0 255.255.255.255
   exit
```

Figure 8-20. Example of Deleting an ACE from Any ACL

Resequencing the ACEs in an ACL

This action reconfigures the starting sequence number for ACEs in an ACL, and resets the numeric interval between sequence numbers for ACEs configured in the ACL.

Syntax: ip access-list resequence < name-str | 1 - 99 | 100 - 199 >

< *starting-seq-#* > < *interval* >

Resets the sequence numbers for all ACEs in the ACL.
*< **starting-seq-#** >* : *Specifies the sequence number for the first ACE in the list. (Default: 10; Range: 1 - 2147483647)*
*< **interval** >* : *Specifies the interval between sequence numbers for the ACEs in the list. (Default: 10; Range: 1 - 2147483647)*

1. To view the current sequence numbering in an ACE, use **show run** or **show access-list < *name-str* | 1 - 99 | 100-199 >**.
2. Use the command syntax (above) to change the sequence numbering.

This example resequences the "My-List" ACL at the bottom of figure 8-20 so that the list begins with line 100 and uses a sequence interval of 100.

```
HP Switch(config)# show run
. . .
ip access-list standard "My-List"
   10 permit 10.10.10.25 0.0.0.0
   15 deny 10.10.10.1 0.0.0.255
   30 deny 10.20.10.1 0.0.0.255
   40 permit 0.0.0.0 255.255.255.255
   exit
. . .
HP Switch(config)# ip access-list resequence My-List 100 100
HP Switch(config)# show run
. . .
ip access-list standard "My-List"
   100 permit 10.10.10.25 0.0.0.0
   200 deny 10.10.10.1 0.0.0.255
   300 deny 10.20.10.1 0.0.0.255
   400 permit 0.0.0.0 255.255.255.255
   exit
```

Figure 8-21. Example of Viewing and Resequencing an ACL

Attaching a Remark to an ACE

A remark is numbered in the same way as an ACE, and uses the same sequence number as the ACE to which it refers. This operation requires that the remark for a given ACE be entered prior to entering the ACE itself.

Syntax: access-list < 1 - 99 | 100 - 199 > remark < *remark-str* >

> *This syntax appends a remark to the end of a numbered ACL and automatically assigns a sequence number to the remark. The next command entry should be the ACE to which the remark belongs. (The new ACE will automatically be numbered with the same sequence number as was used for the preceding remark.)*

Syntax: ip access-list < standard | extended > < *name-str* | 1-99 | 100-199 >
[*seq-#*] remark < *remark-str* >
no < *seq-#* > remark

> *This syntax applies to both named and numbered ACLs. Without an optional sequence number, the remark is appended to the end of the list and automatically assigned a sequence number. When entered with an optional sequence number, the remark is inserted in the list according to the numeric precedence of the sequence number. The* **no** *form of the command deletes the indicated remark, but does not affect the related ACE.*

To associate a remark with a specific ACE, enter the remark first, and then enter the ACE.
> – *Entering a remark without a sequence number and then entering an ACE without a sequence number results in the two entries being automatically paired with the same sequence number and appended to the end of the current ACL.*
> – *Entering a remark with a sequence number and then entering an ACE with the same sequence number results in the two entries being paired together and positioned in the list according to the sequence number they share.*

Note

After a numbered ACL has been created (using **access-list < 1 - 99 | 100 - 199 >**), it can be managed as either a named or numbered ACL. For example, in an existing ACL with a numeric identifier of "115", either of the following command sets adds an ACE denying IPv4 traffic from any source to a host at 10.10.10.100:

```
HP Switch(config)# access-list 115 deny ip host
10.10.10.100

HP Switch(config)# ip access-list extended 115
HP Switch(config-ext-nacl)# deny ip any 10.10.10.100
```

Appending Remarks and Related ACEs to the End of an ACL. To include a remark for an ACE that will be appended to the end of the current ACL, enter the remark first, then enter the related ACE. This results in the

remark and the subsequent ACE having the same sequence number. For example, to add remarks using the "Named-ACL" (**nacl**) context:

```
HP Switch(config)# ip access-list standard My-List
HP Switch(config-std-nacl)# permit host 10.10.10.15
HP Switch(config-std-nacl)# deny 10.10.10.1/24
HP Switch(config-std-nacl)# remark HOST-10.20.10.34
HP Switch(config-std-nacl)# permit host 10.20.10.34
HP Switch(config-std-nacl)# show run
. . .
hostname "HP_switch"
ip access-list standard "My-List"
   10 permit 10.10.10.15 0.0.0.0
   20 deny 10.10.10.1 0.0.0.255
   30 remark "HOST-10.20.10.34"
   30 permit 10.20.10.34 0.0.0.0
   exit
```

The remark is assigned the same number that the immediately following ACE ("30" in this example) is assigned when it is automatically appended to the end of the list. This operation applies where new remarks and ACEs are appended to the end of the ACL and are automatically assigned a sequence number.

Figure 8-22. Example of Appending a Remark and Its Related ACE to the End of an ACL

(You can also perform the operation illustrated in figure 8-22 by using the numbered, **access-list < 1 - 99 | 100 - 199 >** syntax shown at the beginning of this section.)

Inserting Remarks and Related ACEs Within an Existing List. To insert an ACE with a remark within an ACL by specifying a sequence number, insert the numbered remark first, then, using the same sequence number, insert the ACE. (This operation applies only to ACLs accessed using the "Named-ACL" (**nacl**) context.) For example:

```
HP Switch(config-std-nacl)# 15 remark "HOST 10.10.10.21"
HP Switch(config-std-nacl)# 15 permit host 10.10.10.21
HP Switch(config-std-nacl)# show run
ip access-list standard "My-List"
   10 permit 10.10.10.15 0.0.0.0
   15 remark "HOST 10.10.10.21"
   15 permit 10.10.10.21 0.0.0.0
   20 deny 10.10.10.1 0.0.0.255
   30 remark "HOST-10.20.10.34"
   30 permit 10.20.10.34 0.0.0.0
   exit
```

Inserting a remark/ACE pair with the same sequence number requires that the remark (with the desired sequence number) be inserted **before** the ACE with the same number.

Inserting a Remark for an ACE that Already Exists in an ACL. If a sequence number is already assigned to an ACE in a list, you cannot insert a remark by assigning it to the same number. (To configure a remark with the same number as a given ACE, the remark must be configured first.) To assign a remark to the same number as an existing ACE:

1. Delete the ACE.

2. Configure the remark with the number you want assigned to the pair.

3. Re-Enter the deleted ACE with the number used to enter the remark.

Removing a Remark from an Existing ACE. If you want to remove a remark, but want to retain the ACE, do the following:

1. Use the Named ACL context to enter the ACL.

2. Using **show run** or **show access-list** *< list-name >* **config**, note the sequence number and content of the ACE having a remark you want to remove.

3. Delete the ACE.

4. Using the same sequence number, re-enter the ACE.
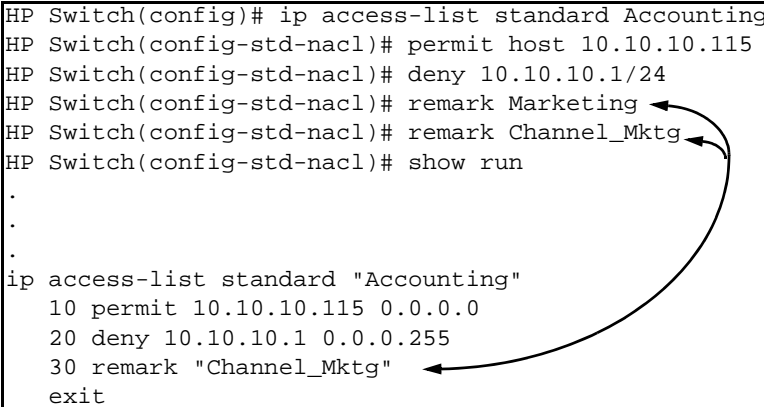
Operating Notes for Remarks

■ The **resequence** command ignores "orphan" remarks that do not have an ACE counterpart with the same sequence number. For example, if:

  • a remark numbered "55" exists in an ACE

  • there is no ACE numbered "55" in the same ACL

  • **resequence** is executed on an ACL

  then the remark retains "55" as its sequence number and will be placed in the renumbered version of the ACL according to that sequence number.

■ Entering an unnumbered remark followed by a numbered ACE, or the reverse, creates an "orphan" remark. The unnumbered entry will be assigned a sequence number that is an increment from the last ACE in the list. The numbered entry will then be placed sequentially in the list according to the sequence number used.

■ Configuring two remarks without either sequence numbers or an intervening, unnumbered ACE results in the second remark over-writing the first.

```
HP Switch(config)# ip access-list standard Accounting
HP Switch(config-std-nacl)# permit host 10.10.10.115
HP Switch(config-std-nacl)# deny 10.10.10.1/24
HP Switch(config-std-nacl)# remark Marketing
HP Switch(config-std-nacl)# remark Channel_Mktg
HP Switch(config-std-nacl)# show run
.
.
.
ip access-list standard "Accounting"
   10 permit 10.10.10.115 0.0.0.0
   20 deny 10.10.10.1 0.0.0.255
   30 remark "Channel_Mktg"
   exit
```

Where multiple remarks are sequentially entered for automatic inclusion at the end of an ACL, each successive remark replaces the previous one until an ACE is configured for automatic inclusion at the end of the list.

Figure 8-23. Example of Overwriting One Remark with Another

Displaying ACL Configuration Data

The show commands in this section apply to both IPv4 and IPv6 ACLs. For information on IPv6 ACL operation, refer to the chapter titled "IPv6 Access Control Lists" in the *IPv6 Configuration Guide* for your switch.

ACL Commands	Function	Page
show access-list	Displays a brief listing of all IPv4 and IPv6 ACLs on the switch.	8-77
show access-list config	Display the type, identifier, and content of all IPv4 and IPv6 ACLs configured in the switch.	8-78
show access-list vlan < *vid* >	List the name and type for each IPv4 and IPv6 ACL application assigned to a particular VLAN on the switch.	8-79
show access-list ports < all | *port-list* >	Lists the IPv4 and IPv6 ACL static port assignment for either all ports and trunks, or for the specified ports and/or trunks.	8-78
show access-list < *acl-name-str* > [config]	Display detailed content for a specific IPv4 or IPv6 ACL. Using the config option displays the ACL in a list format similar to that used to display an ACL in the **show running-config** output.	8-82
show access-list resources	Displays information on the resources currently available in the switch. Refer to the "Monitoring Resources" appendix in the latest *Management and Configuration Guide* for your switch.	—
show port-access web-based clients < *port-list* > detailed show port-access mac-based clients < *port-list* > detailed show port-access authenticator clients < *port-list* > detailed	Shows details for ports in the < *port-list* >.	—
show config show running	**show config** includes configured ACLs and assignments existing in the startup-config file. **show running** includes configured ACLs and assignments existing in the running-config file.	—

Display an ACL Summary

This command lists the configured IPv4 and IPv6 ACLs, regardless of whether they are assigned to any VLANs.

Syntax: show access-list

> *List a summary table of the name, type, and application status of IPv4 and IPv6 ACLs configured on the switch.*

For example:

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

Access Control Lists

 Type  Appl  Name
 ----  ----  ------------------------------------------
   ext  yes  101
   std  yes  55
   ext  yes  Marketing
  ipv6   no  Accounting
  ipv6   no  List-01-Inbound
  ipv6  yes  List-02-Outbound
  ipv6  yes  Test-1
```

IPv4 ACLs

These ACLs exist in the configuration but are not applied to any interfaces and thus do not affect traffic.
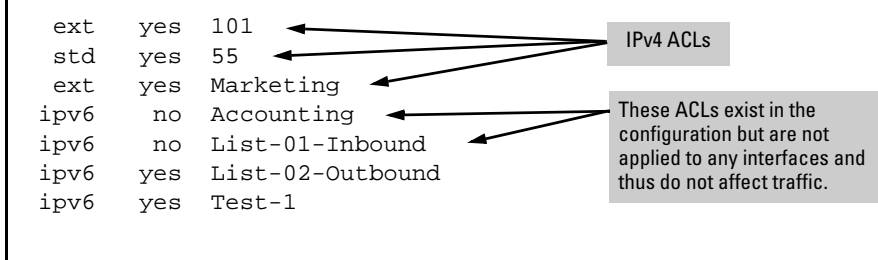
Figure 8-24. Example of a Summary Table of Access lists

Term	Meaning
Type	Shows whether the listed ACL is an IPv4 **std** ACL, an IPv4 **ext** ACL, or an IPv6 ACL.
Appl	Shows whether the listed ACL has been applied to an interface (**yes**/**no**).
Name	Shows the identifier (name or number) assigned to each ACL configured in the switch.

Display the Content of All ACLs on the Switch

This command lists the configuration details for the IPv4 and IPv6 ACLs in the running-config file, regardless of whether any are actually assigned to filter IPv4 traffic on specific VLANs.

Syntax: show access-list config

> *List the configured syntax for all IPv4 and IPv6 ACLs currently configured on the switch.*

Note   Notice that you can use the output from this command for input to an offline text file in which you can edit, add, or delete ACL commands. Refer to "Creating or Editing ACLs Offline" on page 8-86.

This information also appears in the **show running** display. If you executed **write memory** after configuring an ACL, it appears in the **show config** display.

Figure 8-25 shows the ACLs on a switch configured with two IPv6 ACLs named "Accounting" and "List-01-Inbound", and one extended IPv4 ACL named "101":

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

ip access-list extended "101"
   10 permit tcp 10.30.133.27 0.0.0.0 0.0.0.0 255.255.255.255
   20 permit tcp 10.30.155.101 0.0.0.0 0.0.0.0 255.255.255.255
   30 deny ip 10.30.133.1 0.0.0.0 0.0.0.0 255.255.255.255 log
   40 deny ip 10.30.155.1 0.0.0.255 0.0.0.0 255.255.255.255
   exit
ipv6 access-list "Accounting"
    10 permit tcp 2001:db8:0:1af::10:14/128 ::/0 eq 23
    20 permit tcp 2001:db8:0:1af::10:23/128 ::/0 eq 23
    30 deny tcp 2001:db8:0:1af::10/116 ::/0 log
    40 permit ipv6 2001:db8:0:1af::10/116 ::/0
    50 deny ipv6 ::/0 ::/0 log
   exit
ipv6 access-list "List-01-Inbound"
    10 permit icmp fe80::10:60/128 ::/0 dscp 38
    20 permit icmp fe80::10:77/128 ::/0 dscp 38
    30 permit icmp fe80::10:83/128 ::/0 dscp 38
    40 deny icmp ::/0 ::/0 dscp 38
    50 permit ipv6 fe80::10/112 ::/0
    60 deny ipv6 fe80::/64 ::/0
   exit
```

Figure 8-25. Example of an ACL Configured Syntax Listing

Display the VACL Assignments for a VLAN

This command briefly lists the identification and type(s) of IPv4 and IPv6
VACLs currently assigned to a particular VLAN in the running-config file.

Syntax: show access-list vlan < *vid* >

*Lists the current IPv4 and IPv6 ACL assignments to the spec-
ified VLAN (in the running config file).*

Note   This information also appears in the **show running** display. If you execute **write
memory** after configuring an ACL, it also appears in the **show config** display.

For example, the following output shows that all inbound IPv6 traffic and the
inbound and outbound, routed IPv4 traffic are all filtered on VLAN 20.

```
HP Switch(config)# show access-list vlan 20

 Access Lists for VLAN 20

  Inbound Access List: Account-2
  Type: Extended

  Outbound Access List: 101
   Type: Extended

  Ipv6 VACL Access List: Blue-Group

  VACL Access List: None

  Connection Rate Filter Access List: None
```

An extended IPv4 ACL named "Account-2" is assigned to filter routed IPv4 traffic entering the switch on VLAN 20.

An extended IPv4 ACL named "101" is assigned to filter routed IPv4 traffic leaving the switch on VLAN 20.

An IPv6 ACL named "Blue-Group" is assigned to filter IPv6 traffic entering the switch on VLAN 20.

There is no ACL configured to filter all IPv4 traffic entering the switch on VLAN 20.

There is no IPv4 Connection Rate Filter ACL assigned to VLAN 10. Refer to the chapter titled "Virus Throttling (Connection-Rate Filtering)" in the latest *Access Security Guide* for your switch.

Figure 8-26. Example of Listing the ACL Assignments for a VLAN

Display Static Port (and Trunk) ACL Assignments

This command lists the identification and type(s) of current static port ACL assignments to individual switch ports and trunks, as configured in the running-config file. (The switch allows one static port ACL assignment per port.)

Syntax: show access-list ports < all | *port-list* >

Lists the current static port ACL assignments for ports and trunks in the running config file.

Note This information also appears in the **show running** output. If you execute **write memory** after configuring an ACL, it also appears in the **show config** output.

For example, the following output shows IPv4 and IPv6 ACLs configured on various ports and trunks on the switch:

```
HP Switch(config)# show access-list ports all

Access Lists for Port B1

   Inbound Ipv6: List-01-Inbound

Access Lists for Port B12

   Inbound  : 101
   Type     : Extended
   Inbound Ipv6: Accounting

Access Lists for Port Trk2

   Inbound Ipv6: Accounting

Access Lists for Port Trk5

   Inbound  : Marketing
   Type     : Extended
```

An IPv6 ACL is filtering inbound traffic on port B1.

Both an IPv4 ACL and an IPv6 ACL are filtering inbound IPv4 and IPv6 traffic, respectively, on port B12.

An IPv6 ACL is filtering inbound IPv6 traffic on Trunk 2 (Trk2).

An IPv4 ACL is filtering inbound IPv4 traffic on Trunk 5 (Trk5).

Figure 8-27. Example of Listing the ACL Assignments for Ports and Trunks

Displaying the Content of a Specific ACL

This command displays a specific IPv6 or IPv4 ACL configured in the running config file in an easy-to-read tabular format.

Note

This information also appears in the **show running** display. If you execute **write memory** after configuring an ACL, it also appears in the **show config** display.

For information on IPv4 ACL operation, refer to the latest version of the *Access Security Guide* for your switch.

Syntax: show access-list < *identifier* > [config]

> *Display detailed information on the content of a specific ACL configured in the running-config file.*

For example, suppose you configured the following two ACLs in the switch:

Identifier	Type	Desired Action
Accounting	IPv6	• Permit Telnet traffic from these two IPv6 addresses: – 2001:db8:0:1af::10: 14 – 2001:db8:0:1af::10: 24 • Deny Telnet traffic from all other devices in the same subnet. • Permit all other IPv6 traffic from the subnet. • Deny and log any IPv6 traffic from any other source.
List-120	IPv4 Extended	• Permit any TCP traffic from 10.30.133.27 to any destination. • Deny any other IPv4 traffic from 10.30.133.(1-255). • Permit all other IPv4 traffic from any source to any destination.

Use **show access-list < *identifier* >** to inspect a specific IPv6 or IPv4 ACL, as follows:

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

      Access Control Lists

       Name: Accounting
       Type: ipv6
       Applied: Yes  ◄──── Indicates whether the ACL
                           is applied to an interface.

      SEQ  Entry
     ------------------------ Remark Field (Appears if remark configured.) ---------------------
       10   Action: permit                                    Source and Destination
            Remark: Telnet Allowed                            Prefix Lengths
            Src IP: 2001:db8:0:1af::10:14         Prefix Len: 128
            Dst IP: :: ◄── Destination Address    Prefix Len: 0
            Src Port(s):      Dst Port(s): eq 23◄──── TCP Destination Port

       20   Action: permit                        Note: An empty TCP field indicates
            Src IP: 2001:db8:0:1af::10:23         that the TCP port number for that  128
            Dst IP: ::                            field can be any value.            0
            Src Port(s):      Dst Port(s): eq 23

       30   Action: deny (log)
            Src IP: 2001:db8:0:1af::10             Prefix Len: 116
            Dst IP: ::                             Prefix Len: 0
            Src Port(s):      Dst Port(s):

       40   Action: permit
            Src IP: 2001:db8:0:1af::10             Prefix Len: 116
            Dst IP: ::                             Prefix Len: 0
            Src Port(s):      Dst Port(s):
```

Source Address ──► (Src IP line)

TCP Source Port ──► (Src Port(s) line)

Figure 8-28. Example of Listing an IPv6 ACL

```
          HP Switch(config)# show access-list List-120

          Access Control Lists

            Name: List-120
            Type: Extended                  Indicates whether the ACL
            Applied: No                     is applied to an interface.


           SEQ   Entry
                                            Remark Field (Appears if remark configured.).
          ----------------------------------                              -----
            10    Action: permit
                  Remark: Telnet Allowed
   Source Address  Src IP: 10.30.133.27      Mask: 0.0.0.0          Port(s): eq 23
   TCP Source Port  Dst IP: 0.0.0.0          Mask: 255.255.255.255  Port(s):

            20    Action: deny (log)                  Empty field indicates that
                  Src IP: 10.30.133.1       Mask: 0.0.0.255         the destination TCP port
                  Dst IP: 0.0.0.0           Mask: 255.255.255.255   can be any value.

            30    Action: permit
                  Src IP: 0.0.0.0           Mask: 255.255.255.255   Port(s):
                  Dst IP: 0.0.0.0           Mask: 255.255.255.255   Port(s):
```

Figure 8-29. Example of Listing an IPv4 Extended ACL

The **show access-list < identifier > config** command shows the same ACL data as **show access-list < identifier >** but in the format used by the **show < run | config >** commands to list the switch configuration. For example:

```
HP Switch(config)# show access-list List-120 config

ip access-list extended "List-120"
   10 remark "Telnet Allowed"
   10 permit tcp 10.30.133.27 0.0.0.0 eq 23 0.0.0.0 255.255.255.255
   20 deny ip 10.30.133.1 0.0.0.255 0.0.0.0 255.255.255.255 log
   30 permit ip 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255
   exit
```

Figure 8-30. Example of an ACL Listed with the "Config" Option

Table 8-7.   Descriptions of Data Types Included in Show Access-List < *acl-id* > Output

Field	Description
Name	The ACL identifier. Can be a number from 1 to 199, or a name.
Type	Standard or Extended. The former uses only source IPv4 addressing. The latter uses both source and destination IPv4 addressing and also allows TCP or UDP port specifiers.
Applied	"Yes" means the ACL has been applied to a port or VLAN interface. "No" means the ACL exists in the switch configuration, but has not been applied to any interface, and is therefore not in use.
SEQ	The sequential number of the Access Control Entry (ACE) in the specified ACL.
Entry	Lists the content of the ACEs in the selected ACL.
Action	Permit (forward) or deny (drop) a packet when it is compared to the criteria in the applicable ACE and found to match. Includes the optional **log** option, if used, in **deny** or **permit** actions.
Remark	Displays any optional remark text configured for the selected ACE.
IP	**Used for Standard ACLs:** The source IPv4 address to which the configured mask is applied to determine whether there is a match with a packet.
Src IP	**Used for Extended ACLs:** Same as above.
Dst IP	**Used for Extended ACLs:** The source and destination IPv4 addresses to which the corresponding configured masks are applied to determine whether there is a match with a packet.
Mask	The mask configured in an ACE and applied to the corresponding IPv4 address in the ACE to determine whether a packet matches the filtering criteria.
Proto	Used only in extended ACLs to specify the packet protocol type to filter. Must be either IPv4, TCP, or UDP. For TCP protocol selections, includes the **established** option, if configured.
Port(s)	Used only in extended ACLs to show any TCP or UDP operator and port number(s) included in the ACE.

Display All ACLs and Their Assignments in the Routing Switch Startup-Config File and Running-Config File

The **show config** and **show running** commands include in their listings any configured ACLs and any ACL assignments to VLANs. Refer to figure 8-7 (page 8-34) for an example. Remember that **show config** lists the startup-config file and **show running** lists the running-config file.

Creating or Editing ACLs Offline

The section titled "Editing an Existing ACL" on page 8-65 describes how to use the CLI to edit an ACL, and is most applicable in cases where the ACL is short or there is only a minor editing task to perform. The offline method provides a useful alternative to using the CLI for creating or extensively editing a large ACL. This section describes how to:

■   move an existing ACL to a TFTP server

■   use a text (.txt) file format to create a new ACL or edit an existing ACL offline

■   use TFTP to load an offline ACL into the switch's running-config

For longer ACLs that may be difficult or time-consuming to accurately create or edit in the CLI, you can use the offline method described in this section.

N o t e   The **copy** commands can use either **tftp, sftp,** or **xmodem**. So although the following example highlights tftp, bear in mind that **sftp** or **xmodem** can also be used to transfer ACLs to and from the switch.

Creating or Editing an ACL Offline

The Offline Process

1.   Begin by doing one of the following:
    •   To edit one or more existing ACLs, use **copy command-output tftp** to copy the current version of the ACL configuration to a file in your TFTP server. For example, to copy the ACL configuration to a file named **acl-02.txt** in the TFTP directory on a server at 10.28.227.2:

        HP Switch# copy command-output 'show access-list
        config' tftp 10.28.227.2 acl02.txt pc
    •   To create a new ACL, just open a text (.txt) file in the appropriate directory on a TFTP server accessible to the switch.
2.   Use a text editor to create or edit the ACL(s) in the ***.txt** ASCII file format.

If you are replacing an ACL on the switch with a new ACL that uses the same number or name syntax, begin the command file with a **no ip access-list** command to remove the earlier version of the ACL from the switch's running-config file. Otherwise, the switch will append the new ACEs in the ACL you download to the existing ACL. For example, if you planned to use the **copy** command to *replace* ACL "List-120", you would place this command at the beginning of the edited file:

```
no ip access-list extended List-120
```

```
no ip access-list extended List-120
ip access-list extended "List-120"
  10 remark "THIS ACE ALLOWS TELNET"
  10 permit tcp 10.30.133.27 0.0.0.0 eq 23 0.0.0.0 255.255.255.
  20 deny ip 10.30.133.1 0.0.0.255 0.0.0.0 255.255.255.255
  30 deny ip 10.30.155.1 0.0.0.255 0.0.0.0 255.255.255.255
  40 remark "THIS IS THE FINAL ACE IN THE LIST"
  40 permit ip 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255
  exit
```

Removes an existing ACL and replaces it with a new version with the same identity. To append new ACEs to an existing ACL instead of replacing it, you would omit the first line and ensure that the sequence numbering for the new ACEs begin with a number greater than the highest number in the existing list.

Figure 8-31. Example of an Offline ACL File Designed To Replace An Existing ACL

3. Use **copy tftp command-file** to download the file as a list of commands to the switch.

4. The command to assign the ACL to a VLAN was included in the .txt command file. If this is not done in your applications, then the next step is to manually assign the new ACL to the intended VLAN.

5. You can then use the **show run** or **show access-list config** command to inspect the switch configuration to ensure that the ACL was properly downloaded.

```
HP Switch(config)# show run
. . .
ip access-list extended "LIST-20-IN"
   10 remark "THIS ACE APPLIES INBOUND ON VLAN 20"
   10 permit tcp 0.0.0.0 255.255.255.255 10.10.20.98 0.0.0.0 eq 80
   20 permit tcp 0.0.0.0 255.255.255.255 10.10.20.21 0.0.0.0 eq 80
   30 deny tcp 0.0.0.0 255.255.255.255 10.10.20.1 0.0.0.255 eq 80
   40 deny tcp 10.10.20.17 0.0.0.0 10.10.10.100 0.0.0.0 eq 23 log
   50 deny tcp 10.10.20.23 0.0.0.0 10.10.10.100 0.0.0.0 eq 23 log
   60 deny tcp 10.10.20.40 0.0.0.0 10.10.10.100 0.0.0.0 eq 23 log
   70 permit ip 10.10.20.1 0.0.0.255 10.10.10.100 0.0.0.0
   80 remark "VLAN 30 POLICY."
   80 deny ip 10.10.30.1 0.0.0.255 10.10.10.100 0.0.0.0
   90 permit ip 10.10.30.1 0.0.0.255 10.10.10.1 0.0.0.255
   exit
. . .
vlan 20
   name "VLAN20"
   no ip address
   ip access-group "LIST-20-in" in
   exit
```

Note that the comments preceded by " ; " in the .txt source file for this configuration do not appear in the ACL configured in the switch.

As a part of the instruction set included in the .txt file, the ACL is assigned to inbound IPv4 traffic on VLAN 20.

Figure 8-32. Example of Verifying the .txt File Download to the Switch

6. If the configuration appears satisfactory, save it to the startup-config file:

HP Switch(config)# write memory

Enable ACL "Deny" or "Permit" Logging

ACL logging enables the switch to generate a message when IP traffic meets the criteria for a match with an ACE that results in an explicit "deny" or "permit" action. You can use ACL logging to help:

■ Test your network to ensure that your ACL configuration is detecting and denying or permitting the IPv4 traffic you do not want forwarded

■ Receive notification when the switch detects attempts to forward IPv4 traffic you have designed your ACLs to reject (deny) or allow (permit)

The switch sends ACL messages to and optionally to the current console, Telnet, or SSH session. You can use **logging < >** to configure up to six server destinations.

Requirements for Using ACL Logging

■ The switch configuration must include an ACL (1) assigned to a port, trunk, or static VLAN interface and (2) containing an ACE configured with the **deny** or **permit** action and the **log** option.

■ For ACL logging to a server:

• The server must be accessible to the switch and identified in the running configuration.

• The logging facility must be enabled for .

• Debug must be configured to:
– support ACL messages
– send debug messages to the desired debug destination

These requirements are described in more detail under "Enabling ACL Logging on the Switch" on page 8-91.

ACL Logging Operation

When the switch detects a packet match with an ACE and the ACE includes either the **deny** or **permit** action, and the optional **log** parameter, an ACL log message is sent to the designated debug destination.

The first time a packet matches an ACE with **deny** or **permit** and **log** configured, the message is sent immediately to the destination and the switch starts a wait-period of approximately five minutes. (The exact duration of the period depends on how the packets are internally routed.) At the end of the collection period, the switch sends a single-line summary of any additional "deny" or "permit" matches for that ACE (and any other "deny" or "permit" ACEs for which the switch detected a match).

If no further log messages are generated in the wait-period, the switch suspends the timer and resets itself to send a message as soon as a new "deny" or "permit" match occurs. If subsequent packets matching the already logged ACL entries are detected, then a new logged event will be generated that summarizes the number of packets that matched each specific entry (with the time period). The data in the message includes the information illustrated in Figure 8-33.

```
Feb 1 10:04:45 10.10.20.1 ACL:
ACL 02/01/07 10:04:45 List NO-TELNET, seq#10 denied
tcp 10.10.10.3(1612)->10.10.20.2(23) on vlan 1, port A7


Feb 1 10:04:45 10.10.20.1 ACL:
ACL 02/01/07 10:04:45 : ACL NO-TELNET seq#10 denied 6 packets
```

Example report of the first **deny** event detected by the switch for this ACE.

Example of subsequent **deny** events detected by the switch for the same ACE.
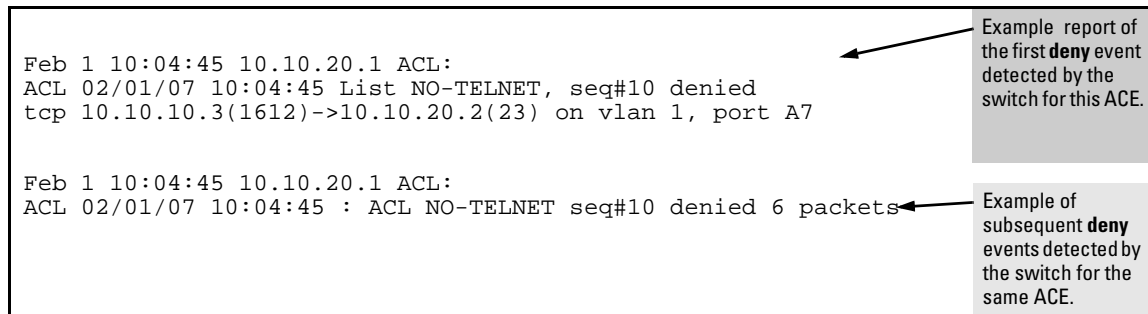
Figure 8-33. Content of a Message Generated by an ACL-Deny Action

Enabling ACL Logging on the Switch

1. If you are using a server, use the **logging < *ip-addr* >** command to configure the server IPv4 address(es). Ensure that the switch can access any server(s) you specify.

2. Use **logging facility** to enable the logging for operation.

3. Use the **debug destination** command to configure one or more log destinations. (Destination options include **logging** and **session**. For more information on debug, refer to "Debug and Messaging Operation" in appendix C, "Troubleshooting", in the *Management and Configuration Guide* for your switch.)

4. Use **debug acl** or **debug all** to configure the debug operation to include ACL messages.

5. Configure one or more ACLs with the **deny** or **permit** action and the **log** option.

Configuring the Logging Timer

By default, the wait period for logging "deny" or "permit" matches (described above in "ACL Logging Operation") is five minutes (300 seconds). You can manually set the wait period timer to an interval between 30 and 300 seconds, using the **access-list** command from the **config** context. This setting is stored in the switch configuration.

Syntax: access-list logtimer <default | <30-300>>

> *From* **config** *context: This command sets the wait period timer for logging "deny" or "permit" messages to the server or other destination device(s). The first time a packet matches an ACE with* **deny** *or* **permit** *and* **log** *configured, the message is sent immediately to the destination and the switch starts a wait period of approximately five minutes (default value). (The exact duration of the period depends on how the packets are internally routed.) At the end of the wait period, the switch sends a single-line summary of any additional "deny" or "permit" matches for that ACE (and any other "deny" or "permit" ACEs for which the switch detected a match). If no further log messages are generated in the wait period, the switch suspends the timer and resets itself to send a message as soon as a new "deny" or "permit" match occurs.*
>
> - **default** — *sets the wait period timer to 300 seconds.*
> - **<30-300>** — *sets the wait period timer to the specified number of seconds.*

Monitoring Static ACL Performance

ACL statistics counters provide a means for monitoring ACL performance by using counters to display the current number of matches the switch has detected for each ACE in an ACL assigned to a switch interface. This can help, for example, to determine whether a particular traffic type is being filtered by the intended ACE in an assigned list, or if traffic from a particular device or network is being filtered as intended.

Syntax: < show | clear > statistics
aclv4 < *acl-name-str* > port < *port-#* >
aclv4 < *acl-name-str* > vlan < *vid* > <vlan >

aclv6 < *acl-name-str* > port < *port-#* >
aclv6 < *acl-name-str* > vlan < *vid* > <vlan >

Displays the current match (hit) count per ACE for the speci-fied IPv6 or IPv4 static ACL assignment on a specific interface.

show: *Displays the current match (hit) count per ACE for the specified IPv6 or IPv4 static ACL assignment on a specific interface.*

clear: *Resets ACE hit counters to zero for the specified IPv6 or IPv4 static ACL assignment on a specific interface.*

Total: *This column lists the running total of the matches the switch has detected for the ACEs in an applied ACL since the ACL's counters were last reset to 0 (zero).*

```
HP Switch# show statistics aclv6 IPV6-ACL vlan 20 vlan

 HitCounts for ACL IPV6-ACL

  Total

(     12)     10 permit icmp ::/0 fe80::20:2/128
(      6)     20 deny tcp ::/0 fe80::20:2/128 eq 23 log
(     41)     30 permit ipv6 ::/0 ::/0

HP Switch# show statistics aclv4 102 vlan 20 vlan

 HitCounts for ACL 102

  Total

(      4)     10 permit icmp 10.10.20.3 0.0.0.0 10.10.20.2 0.0.0.0 8
(      8)     20 deny icmp 0.0.0.0 255.255.255.255 10.10.20.2 0.0.0.0 8
(      2)     30 permit tcp 10.10.20.3 0.0.0.255 10.10.20.2 0.0.0.255 eq 23
(      2)     55 deny tcp 0.0.0.0 255.255.255.255 10.10.20.2 0.0.0.0 8
(    125)     60 permit ip 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255
```

Figure 8-34. Example of IPv6 and IPv4 ACL Statistics

ACE Counter Operation: *For a given ACE in an assigned ACL, the counter increments by 1 each time the switch detects a packet that matches the criteria in that ACE, and maintains a running total of the matches since the last counter reset.*

Note: *This ACL monitoring feature does not include hits on the "implicit deny" that is included at the end of all ACLs.*

Resetting ACE Hit Counters to Zero:

- *Using the* **clear statistics** *command (page 8-93).*
- *Removing an ACL from an interface zeros the ACL's ACE counters for that interface only.*
- *For a given ACL, either of the following actions clear the ACE counters to zero for all interfaces to which the ACL is assigned.*
  - *adding or removing a permit or deny ACE in the ACL*
  - *rebooting the switch*

Example of ACL Performance Monitoring

Figure 8-35 shows a sample of performance monitoring output for an IPv6 ACL assigned as a VACL.

```
HP Switch# show statistics aclv6 V6-02 vlan 20 vlan

 HitCounts for ACL V6-02

  Total

(      5)     10 permit icmp ::/0 fe80::20:2/128
(      4)     20 permit icmp ::/0 fe80::20:3/128
(    136)     30 permit tcp fe80::20:1/128 ::/0 eq 23
(      2)     40 deny icmp ::/0 fe80::20:1/128
(     10)     50 deny tcp ::/0 ::/0 eq 23
(      8)     60 deny icmp ::/0 ::/0
(    155)     70 permit ipv6 ::/0 ::/0
```

Figure 8-35. Example of IPv6 ACL Performance Monitoring Output

Figure 8-36 shows a sample of performance monitoring output for an IPv4 ACL assigned as a VACL.

```
HP Switch# show statistics aclv4 102 vlan 20 vlan

 HitCounts for ACL 102

  Total

(      1)      10 permit icmp 10.10.20.3 0.0.0.0 10.10.20.2 0.0.0.0 8
(      2)      20 deny icmp 10.10.20.3 0.0.0.0 10.10.20.1 0.0.0.0 8 log

(      2)      30 deny icmp 10.10.20.2 0.0.0.0 10.10.20.3 0.0.0.0 8 log

(      1)      40 deny icmp 10.10.20.2 0.0.0.0 10.10.20.1 0.0.0.0 8 log

(     10)      50 deny tcp 10.10.20.2 0.0.0.255 10.10.20.3 0.0.0.255 eq 23 log
(     27)      60 permit ip 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255
```

Figure 8-36. Example of IPv4 ACL Performance Monitoring Output

Example of Resetting ACE Hit Counters to Zero

The following example uses the counter activity in figure 8-35 (page 8-94) to demonstrate using **clear statistics** to reset the counters to zero.

```
HP Switch# show statistics aclv6 V6-02 vlan 20 vlan

 HitCounts for ACL V6-02

  Total

(       5)     10 permit icmp ::/0 fe80::20:2/128 128
(       4)     20 permit icmp ::/0 fe80::20:3/128 128
(     136)     30 permit tcp fe80::20:1/128 ::/0 eq 23
(       2)     40 deny icmp ::/0 fe80::20:1/128 128
(      10)     50 deny tcp ::/0 ::/0 eq 23
(       8)     60 deny icmp ::/0 ::/0 133
(     155)     70 permit ipv6 ::/0 ::/0
HP Switch# clear statistics aclv6 V6-02 vlan 20 vlan
HP Switch# show statistics aclv6 V6-02 vlan 20 vlan

 HitCounts for ACL V6-02

  Total


(       0)     10 permit icmp ::/0 fe80::20:2/128 128
(       0)     20 permit icmp ::/0 fe80::20:3/128 128
(       0)     30 permit tcp fe80::20:1/128 ::/0 eq 23
(       0)     40 deny icmp ::/0 fe80::20:1/128 128
(       0)     50 deny tcp ::/0 ::/0 eq 23
(       0)     60 deny icmp ::/0 ::/0 133
(       0)     70 permit ipv6 ::/0 ::/0
```

Figure 8-37. Example of IPv6 ACL Performance Monitoring Output

IPv6 Counter Operation with Multiple Interface Assignments

Note

The examples of counters in this section use small values to help illustrate counter operation. The counters in real-time network applications are generally much more active and show higher values.

Where the same IPv6 ACL is assigned to multiple interfaces, the switch maintains a separate instance of each ACE counter in the ACL. When there is a match with traffic on one of the ACL's assigned interfaces, only the affected ACE counters for that interface are incremented. Other instances of the same ACL applied to other interfaces are not affected.

For example, suppose that:

■ An ACL named "V6-01" is configured as shown in figure to block Telnet access to a workstation at FE80::20:2, which is connected to a port belonging to VLAN 20.

■ The ACL is assigned as a PACL (port ACL) on port B2, which is also a member of VLAN 20:

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

ipv6 access-list "V6-01"
    10 permit icmp ::/0 fe80::20:2/128 128
    20 deny tcp ::/0 fe80::20:2/128 eq 23 log
    30 permit ipv6 ::/0 ::/0
   exit                                      Assigns the ACL to port B2.

HP Switch(config)# int b2 ipv access-group V6-01 in
```
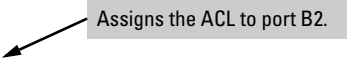
Figure 8-38. ACL "V6-01" and Command for PACL Assignment on Port B2
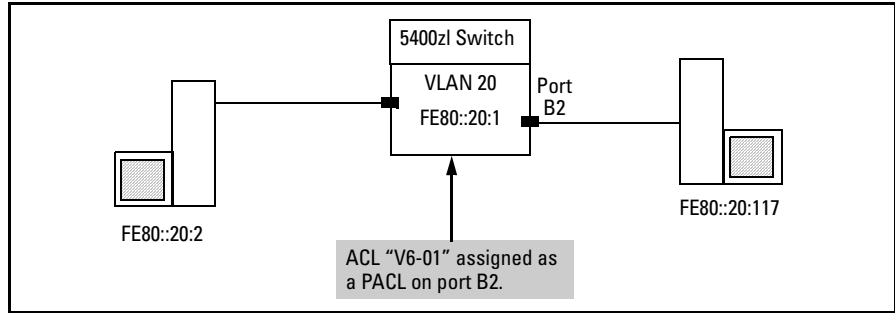
Figure 8-39. Application to Filter Traffic Inbound on Port B2

Using the topology in figure 8-39, a workstation at FE80::20:117 on port B2 attempting to ping and Telnet to the workstation at FE80::20:2 is filtered through the PACL instance of the "V6-01" ACL assigned to port B2, resulting in the following:

```
HP Switch# ping6 fe80::20:2%vlan20
fe80:0000:0000:0000:0000:0000:0020:0002 is alive, time = 5 ms
HP Switch# telnet fe80::20:2%vlan20
Telnet failed: Connection timed out.
HP Switch#
```

Figure 8-40. Ping and Telnet from FE80::20:117 to FE80::20:2 Filtered by the Assignment of "V6-01" as a PACL on Port B2

Figure 8-41. Resulting ACE Hits on ACL "V6-01"

General ACL Operating Notes

ACLs do not provide DNS hostname support. ACLs cannot be configured to screen hostname IPv4 traffic between the switch and a DNS.

ACLs Do Not Affect Serial Port Access. ACLs do not apply to the switch's serial port.

ACL Screening of IPv4 Traffic Generated by the Switch. t ACLs applied on the switch screen IPv4 traffic when other devices generate it. Similarly, ACL applications can screen responses from other devices to unscreened IPv4 traffic the switch generates.

ACL Logging.

- The ACL logging feature generates a message only when packets are explicitly denied or permitted as the result of a match, and not when implicitly denied. To help test ACL logging, configure the last entry in an ACL as an explicit **deny** or **permit** statement with a **log** statement included, and apply the ACL to an appropriate VLAN.

- A detailed event will be logged for the first packet that matches a "deny" or "permit" ACL logged entry with the appropriate action specified.

  Subsequent packets matching ACL logged entries will generate a new event that summarizes the number of packets that matched each specific entry (with the time period). See Figure 8-33 for an example.

- Logging enables you to selectively test specific devices or groups. However, *excessive logging can affect switch performance*. For this reason, HP recommends that you remove the logging option from ACEs for which you do not have a present need. Also, avoid configuring logging where it does not serve an immediate purpose. (Note that ACL logging is not designed to function as an accounting method.) See also "Apparent Failure To Log All 'Deny' or 'Permit' Matches" in the section titled "ACL Problems", found in appendix C, "Troubleshooting" of the *Management and Configuration Guide* for your switch.

- When configuring logging, you can reduce excessive resource use by configuring the appropriate ACEs to match with specific hosts instead of entire subnets. (For more on resource usage, refer to "Monitoring Shared Resources" on page 8-100.)

Minimum Number of ACEs in an ACL. Any ACL must include at least one ACE to enable IP traffic screening. A numbered ACL cannot be created without at least one ACE. A named ACL can be created "empty"; that is, without any ACEs. However in an empty ACL applied to an interface, the Implicit Deny function does not operate, and the ACL has no effect on traffic.

Monitoring Shared Resources. Applied ACLs share internal switch resources with several other features. The switch provides ample resources for all features. However, if the internal resources become fully subscribed, additional ACLs cannot be applied until the necessary resources are released from other applications. For information on determining current resource availability and usage, refer to appendix E, "Monitoring Resources" in the *Management and Configuration Guide* for your switch.

Replacing or Adding To an Active ACL Policy. If you assign an ACL to an interface and subsequently add or replace ACEs in that ACL, each new ACE becomes active when you enter it. If the ACL is configured on multiple interfaces when the change occurs, then the switch resources must accommodate all applications of the ACL. If there are insufficient resources to accommodate one of several ACL applications affected by the change, then the change is not applied to any of the interfaces and the previous version of the ACL remains in effect. Refer to "Monitoring Shared Resources", above.

"Strict" TCP and UDP. When the ACL configuration includes TCP or UDP options, the switch operates in "strict" TCP and UDP mode for increased control. In this case, the switch compares all TCP and UDP packets against the ACLs. (In the HP 9300m and 9404sl Routing Switches, the Strict TCP and Strict UDP modes are optional and must be specifically invoked.)

Traffic/Security Filters and Monitors

Introduction

Feature	Default	Menu	CLI	WebAgent
configure source-port filters	none	n/a	page 9-17	n/a
configure multicast filters	none	n/a	page 9-17	n/a
display filter data	n/a	n/a	page 9-19	n/a

You can enhance in-band security and improve control over access to network resources by configuring static filters to forward (the default action) or drop unwanted traffic. That is, you can configure a traffic filter to either forward or drop all network traffic moving to outbound (destination) ports and trunks (if any) on the switch.

Filter Limits

These limitations apply:

■   Source-port filters: up to 78

■   Multicast filters: up to 16 with 1024 or fewer VLANs configured. Up to 8 with more than 1024 VLANs configured.

Using Port Trunks with Filters

The switch manages a port trunk as a single source or destination for source-port filtering. If you configure a port for filtering before adding it to a port trunk, the port retains the filter configuration, but suspends the filtering action while a member of the trunk. If you want a trunk to perform filtering, first configure the trunk, then configure the trunk for filtering. Refer to "Configuring a Filter on a Port Trunk" on page 9-15.

Filter Types and Operation

Table 9-1. **Filter Types and Criteria**

Static Filter Type	Selection Criteria
Source-Port	Inbound traffic from a designated, physical source-port will be forwarded or dropped on a per-port (destination) basis.
Multicast	Inbound traffic having a specified multicast MAC address will be forwarded to outbound ports (the default) or dropped on a per-port (destination) basis.

Using Source-Port Filters

This filter type enables the switch to forward or drop traffic from *all* end nodes on the indicated source-port to specific destination ports.

Configuring a source-port filter to drop traffic received on port 1 with an outbound destination of port 2 means that End Nodes A, B, and C cannot send traffic to the server. To block traffic in the opposite direction, you would also configure a source-port filter to drop traffic received on port 2 with an outbound destination of port 1.
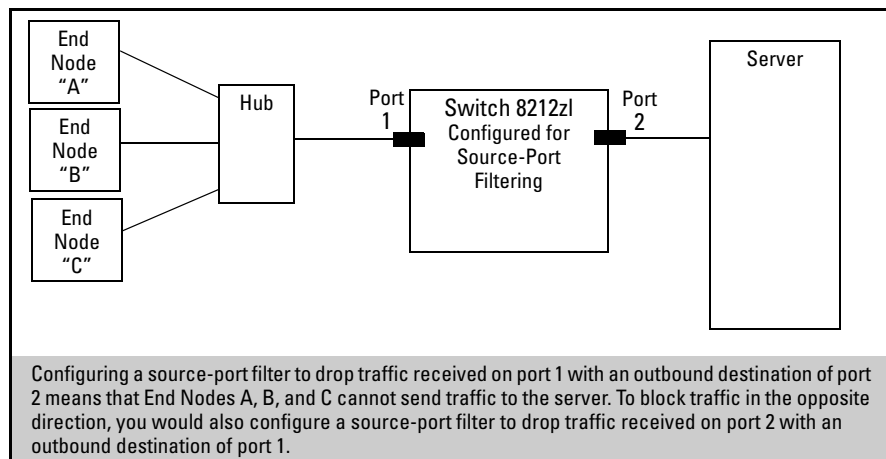
Figure 9-1. Example of a Source-Port Filter Application

Operating Rules for Source-Port Filters

■ You can configure one source-port filter for each physical port and port trunk on the switch. (Refer to the **filter** command on page 9-13.)

■ You can include all destination ports and trunks in the switch on a single source-port filter.

■   Each source-port filter includes:

- One source port or port trunk (**trk1**, **trk2**, ...**trk***n*)

- A set of destination ports and/or port trunks that includes all untrunked LAN ports and port trunks on the switch

- An action (forward or drop) for each destination port or port trunk

When you create a source-port filter, the switch automatically sets the filter to forward traffic from the designated source to all destinations for which you do not specifically configure a "drop" action. Thus, it is not necessary to configure a source-port filter for traffic you want the switch to forward unless the filter was previously configured to drop the desired traffic.

■   When you create a source port filter, all ports and port trunks (if any) on the switch appear as destinations on the list for that filter, even if routing is disabled and separate VLANs and/or subnets exist. Where traffic would normally be allowed between ports and/or trunks, the switch automatically forwards traffic to the outbound ports and/or trunks you do not specifically configure to drop traffic. (Destination ports that comprise a trunk are listed collectively by the trunk name—such as **Trk1**— instead of by individual port name.)

■   Packets allowed for forwarding by a source-port filter are subject to the same operation as inbound packets on a port that is not configured for source-port filtering.

■   With multiple IP addresses configured on a VLAN, and routing enabled on the switch, a single port or trunk can be both the source and destination of packets moving between subnets in that same VLAN. In this case, you can prevent the traffic of one subnet from being routed to another subnet of the same port by configuring the port or trunk as both the source and destination for traffic to drop.

Example

If you wanted to prevent server "A" from receiving traffic sent by workstation "X", but do not want to prevent any other servers or end nodes from receiving traffic from workstation "X", you would configure a filter to drop traffic from port 5 to port 7. The resulting filter would drop traffic from port 5 to port 7, but would forward all other traffic from any source port to any destination port. (Refer to figures 9-2 and 9-3.)

Figure 9-2. Example of a Filter Blocking Traffic only from Port 5 to Server "A"

Figure 9-3. The Filter for the Actions Shown in Figure 9-2

Named Source-Port Filters

You can specify named source-port filters that may be used on multiple ports and port trunks. A port or port trunk can only have one source-port filter, but by using this capability you can define a source-port filter once and apply it to multiple ports and port trunks. This can make it easier to configure and manage source-port filters on your switch. The commands to define, configure, apply, and display the status of named source-port filters are described below.

Operating Rules for Named Source-Port Filters

■ A port or port trunk may only have one source-port filter, named or not named.

■ A named source-port filter can be applied to multiple ports or port trunks.

■ Once a named source-port filter is defined, subsequent changes only modify its action, they don't replace it.

■ To change the named source-port filter used on a port or port trunk, the current filter must first be removed, using the **no filter source-port named-filter <filter-name >** command.

■ A named source-port filter can only be deleted when it is not applied to any ports.

Defining and Configuring Named Source-Port Filters

The named source-port filter command operates from the global configuration level.

Syntax: [no] filter source-port named-filter *<filter-name>*

> *Defines or deletes a named source-port filter. The <filter-name> may contain a maximum of 20 alpha-numeric characters (longer names may be specified, but they are not displayed.) A filter-name cannot be a valid port or port trunk name.*
>
> *The maximum number of named source-port filters that can be used is equal to the number of ports on a switch.*
>
> *A named source-port filter can only be removed if it is not in use (use the* **show filter source-port** *command to check the status). Named source-port filters are not automatically deleted when they are no longer used.*
>
> *Use the* **no** *option to delete an unused named source-port filter.*

Syntax: filter source-port named-filter <*filter-name* > drop < *destination-port-list* >

> *Configures the named source-port filter to drop traffic having a destination on the ports and/or port trunks in the <*
> *destination-port-list >. Can be followed by the* **forward** *option if you have other destination ports or port trunks previously set to* **drop** *that you want to change to* **forward**. *For example:*
> *filter source-port named-filter <filter-name > drop < destina-tion-port-list > forward < destination-port-list>*
> *The* **destination-port-list** *may contain ports, port trunks, and ranges (for example 3-7 or trk4-trk9) separated by commas.*

Syntax: filter source-port named-filter <*filter-name* > forward
< *destination-port-list* >

> *Configures the named source-port filter to forward traffic having a destination on the ports and/or port trunks in the <***destination-port-list***>. Since "forward" is the default state for destinations in a filter, this command is useful when destinations in an existing filter are configured for "drop" and you want to change them to "forward". Can be followed by the* **drop** *option if you have other destination ports set to* **forward** *that you want to change to* **drop**. *For example:*
> filter source-port named-filter <*filter-name* > forward <
> *destination-port-list* > drop < *destination-port-list* >

A named source-port filter must first be defined and configured before it can be applied. In the following example two named source-port filters are defined, **web-only** and **accounting**.

```
HP Switch(config)# filter source-port named-filter web-
only

HP Switch(config)# filter source-port named-filter
accounting
```

By default, these two named source-port filters forward traffic to all ports and port trunks.

To configure a named source-port filter to prevent inbound traffic from being forwarded to specific destination switch ports or port trunks, the **drop** option is used. For example, on a 26-port switch, to configure the named source-port filter **web-only** to drop any traffic except that for destination ports 1 and 2, the following command would be used:

```
HP Switch(config)# filter source-port named-filter web-
only drop 3-26
```

A named source-port filter can be defined and configured in a single command by adding the **drop** option, followed by the required destination-port-list.

Viewing a Named Source-Port Filter

You can list all source-port filters configured in the switch, both named and unnamed, and their action using the **show** command below.

Syntax: show filter source-port

> *Displays a listing of configured source-port filters, where each filter entry includes a Filter Name, Port List, and Action:*
>
> > ***Filter Name**: The filter-name used when a named source-port filter is defined. Non-named source-port filters are automatically assigned the port or port trunk number of the source port.*
> >
> > ***Port List**: Lists the port and port trunk destinations using the filter. Named source-port filters that are not in use display **NOT USED**.*
> >
> > ***Action**: Lists the ports and port trunks dropped by the filter. If a named source-port filter has been defined but not configured, this field is blank.*

[index]   *For the supplied index (IDX) displays the action taken (Drop or Forward) for each destination port on the switch.*

Using Named Source-Port Filters

A company wants to manage traffic to the Internet and its accounting server on a 26-port switch. Their network is pictured in Figure 9-4. Switch port 1 connects to a router that provides connectivity to a WAN and the Internet. Switch port 7 connects to the accounting server. Two workstations in accounting are connected to switch ports 10 and 11.

Network Design

1. Accounting Workstations may only send traffic to the Accounting Server.
2. No Internet traffic may be sent to the Accounting Server or Workstations.
3 All other switch ports may only send traffic to Port 1.
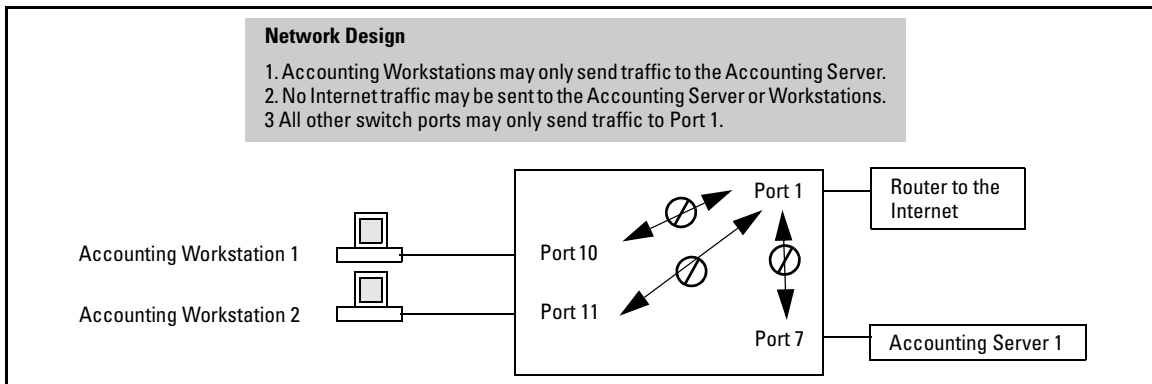
Figure 9-4. Network Configuration for Named Source-Port Filters Example

Defining and Configuring Example Named Source-Port Filters. While named source-port filters may be defined and configured in two steps, this is not necessary. Here we define and configure each of the named source-port filters for our example network in a single step.

```
HP Switch(config)# filter source-port named-filter web-only drop 2-26
HP Switch(config)# filter source-port named-filter accounting drop 1-6,8,9,12-26
HP Switch(config)# filter source-port named-filter no-incoming-web drop 7,10,11

HP Switch(config)# show filter source-port

Traffic/Security Filters

Filter Name          | Port List           | Action
 ------------------- + ------------------- + -------------------
web-only             | NOT USED            | drop 2-26
accounting           | NOT USED            | drop 1-6,8-9,12-26
no-incoming-web      | NOT USED            | drop 7,10-11

HP Switch(config)#
```

Ports and port trunks using the filter. When **NOT USED** is displayed the named source-port filter may be deleted.

Lists the ports and port trunks dropped by the filter. Ports and port trunks not shown are forwarded by the filter.

To remove a port or port trunk from the list, update the named source-port filter definition using the **forward** option.

Figure 9-5. Applying Example Named Source-Port Filters

Once the named source-port filters have been defined and configured we now apply them to the switch ports.

```
HP Switch(config)# filter source-port 2-6,8,9,12-26 named-filter web-only
HP Switch(config)# filter source-port 7,10,11 named-filter accounting
HP Switch(config)# filter source-port 1 named-filter no-incoming-web
HP Switch(config)#
```

Figure 9-6. Source Port Filters Applied to Switch Ports

The **show filter** command shows what ports have filters applied.

```
HP Switch(config)# show filter

Traffic/Security Filters

 IDX Filter Type   | Value
 --- ------------- + --------------------
 1    Source Port  | 2
 2    Source Port  | 3
 3    Source Port  | 4
 4    Source Port  | 5
 5    Source Port  | 6
 6    Source Port  | 8
 7    Source Port  | 9
 8    Source Port  | 12

 20   Source Port  | 24
 21   Source Port  | 25
 22   Source Port  | 26
 23   Source Port  | 7
 24   Source Port  | 10
 25   Source Port  | 11
 26   Source Port  | 1
```

Indicates the port number or port-trunk name of the source port or trunk assigned to the filter.

An automatically assigned index number used to identify the filter for a detailed information listing. A filter retains its assigned IDX number for as long as the filter exists in the switch. The switch assigns the lowest available IDX number to a new filter. This can result in a newer filter having a lower IDX number than an older filter if a previous (source-port or named source-port) filter deletion created a gap in the filter listing.

Figure 9-7. Example of the show filter Command

Using the **IDX** value in the **show filter** command, we can see how traffic is filtered on a specific port (**Value**).The two outputs below show a non-accounting and an accounting switch port.

```
HP Switch(config)# show filter 4        HP Switch(config)# show filter 24

 Traffic/Security Filters                Traffic/Security Filters

  Filter Type : Source Port               Filter Type : Source Port
  Source Port : 5                         Source Port : 10

  Dest   Port Type     | Action          Dest   Port Type      | Action
  -------- --------- + -------           -------- --------- + -----------
  1       10/100TX    | Forward          1        10/100TX    | Drop
  2       10/100TX    | Drop             2        10/100TX    | Drop
  3       10/100TX    | Drop             3        10/100TX    | Drop
  4       10/100TX    | Drop             4        10/100TX    | Drop
  5       10/100TX    | Drop             5        10/100TX    | Drop
  6       10/100TX    | Drop             6        10/100TX    | Drop
  7       10/100TX    | Drop             7        10/100TX    | Forward
  8       10/100TX    | Drop             8        10/100TX    | Drop
  9       10/100TX    | Drop             9        10/100TX    | Drop
  10      10/100TX    | Drop             10       10/100TX    | Drop
  11      10/100TX    | Drop             11       10/100TX    | Drop
  12      10/100TX    | Drop             12       10/100TX    | Drop
   .          .          .                .          .           .
```

Figure 9-8. Example Showing Traffic Filtered on Specific Ports

The same command, using IDX 26, shows how traffic from the Internet is handled.

```
HP Switch(config)# show filter 26

 Traffic/Security Filters

  Filter Type : Source Port
  Source Port : 1

  Dest    Port Type      | Action
  --------- --------- + -----------------------
  1         10/100TX      | Forward
  2         10/100TX      | Forward
  3         10/100TX      | Forward
  4         10/100TX      | Forward
  5         10/100TX      | Forward
  6         10/100TX      | Forward
  7         10/100TX      | Drop
  8         10/100TX      | Forward
  9         10/100TX      | Forward
  10        10/100TX      | Drop
  11        10/100TX      | Drop
  12        10/100TX      | Forward
   .         .        .
```

Figure 9-9. Example of Source Port Filtering with Internet Traffic

As the company grows, more resources are required in accounting. Two additional accounting workstations are added and attached to ports 12 and 13. A second server is added attached to port8.

Network Design

1. Accounting Workstations may only send traffic to the Accounting Server.
2. No Internet traffic may be sent to the Accounting Server or Workstations.
3 All other switch ports may only send traffic to Port 1.

Accounting Workstation 1 — Port 10 — Port 1 — Router to the Internet
Accounting Workstation 2 — Port 11
Accounting Workstation 3 — Port 12 — Port 7 — Accounting Server 1
Accounting Workstation 4 — Port 13 — Port 8 — Accounting Server 2

Figure 9-10. Expanded Network Configuration for Named Source-Port Filters Example

The following revisions to the named source-port filter definitions maintain the desired network traffic management, as shown in the **Action** column of the **show** command.

```
HP Switch(config)# filter source-port named-filter accounting forward 8,12,13
HP Switch(config)# filter source-port named-filter no-incoming-web drop 8,12,13
HP Switch(config)#
HP Switch(config)# show filter source-port

 Traffic/Security Filters

 Filter Name           | Port List              | Action
 -------------------   + -------------------    + ----------------------
 web-only              | 2-6,8-9,12-26          | drop 2-26
 accounting            | 7,10-11                | drop 1-6,9,14-26
 no-incoming-web       | 1                      | drop 7-8,10-13

HP Switch(config)#
```

Figure 9-11. Example Showing Network Traffic Management with Source Port Filters

We next apply the updated named source-port filters to the appropriate switch ports. As a port can only have one source-port filter (named or not named), before applying the new named source-port filters we first remove the existing source-port filters on the port.

```
HP Switch(config)# no filter source-port 8,12,13
HP Switch(config)# filter source-port 8,12,13 named-filter accounting
HP Switch(config)#
```

The named source-port filters now manage traffic on the switch ports as shown below, using the **show filter source-port** command.

```
HP Switch(config)# show filter source-port

 Traffic/Security Filters

 Filter Name           | Port List            | Action
 -------------------- + -------------------- + -------------------------
 web-only              | 2-6,9,14-26          | drop 2-26
 accounting            | 7-8,10-13            | drop 1-6,9,14-26
 no-incoming-web       | 1                    | drop 7-8,10-13

HP Switch(config)#
```

Figure 9-12. Named Source-Port Filters Managing Traffic Configuring a Source-Port Traffic Filter

Configuring a Source-Port Filter

The source-port filter command is executed from the gloabl configuration level.

Syntax: [no] filter

[source-port < *port-number* | *trunk-name*>]

Specifies one inbound port or trunk. Traffic received inbound on this interface from other devices will be filtered. The **no** *form of the command deletes the source-port filter for <* port-number *> and returns the destination ports for that filter to the* **Forward** *action. (Default: Forward on all ports.)*

Note: If multiple VLANs are configured, the source-port and the destination port(s) must be in the same VLAN unless routing is enabled. Similarly, if a VLAN containing both the source and destination is multi-netted, the source and destination ports and/or trunks must be in the same subnet unless routing is enabled.

[drop] < *destination-port-list* > [forward < *port-list* >]

> *Configures the filter to drop traffic for the ports and/or trunks in the designated* **< destination-port-list >**. *Can be followed by* **forward < destination-port-list >** *if you have other destination ports set to* **drop** *that you want to change to* **forward**. *If no drop or forward action is specified, the switch automatically creates a filter with a* **forward** *action from the designated source port (or trunk) to all destination ports (or trunks) on the switch.*

[forward] < *port-list* >

> *Configures the filter to forward traffic for the ports and/or trunks in the designated* **< destination-port-list >**. *Because* **forward** *is the default state for destinations in a filter, this command is useful when destinations in an existing filter are configured for* **drop** *and you want to change them to* **forward**. *Can be followed by* **drop < destination-port-list >** *if you have other destination ports set to* **forward** *that you want to change to* **drop**. *If no drop or forward action is specified, the switch automatically creates a filter with a forward action from the designated source port (or trunk) to all destination ports (or trunks) on the switch.*

Example of Creating a Source-Port Filter

For example, assume that you want to create a source-port filter that drops all traffic received on port 5 with a destination of port trunk 1 (**Trk1**) and any port in the range of port 10 to port 15. To create this filter you would execute this command:

```
HP Switch(config)# filter source-port 5 drop trk1,10-15
```

Later, suppose you wanted to shift the destination port range for this filter up by two ports; that is, to have the filter drop all traffic received on port 5 with a destination of any port in the range of port 12 to port 17. (The **Trk1** destination is already configured in the filter and can remain as-is.) With one command you can restore forwarding to ports 10 and 11 while adding ports 16 and 17 to the "drop" list:

```
HP Switch(config)# filter source-port 5 forward 10-11 drop
16-17
```

Configuring a Filter on a Port Trunk

This operation uses the same command as is used for configuring a filter on an individual port. However, the configuration process requires two steps:

1. Configure the port trunk.

2. Configure a filter on the port trunk by using the trunk name (**trk1**, **trk2**, ...**trk6**) instead of a port name.

For example, to create a filter on port trunk 1 to drop traffic received inbound for trunk 2 and ports 10-15:

```
HP Switch(config)# filter source-port trk1 drop trk2,10-15
```

Note that if you first configure a filter on a port and then later add the port to a trunk, the port remains configured for filtering *but the filtering action will be suspended while the port is a member of the trunk*. That is, the trunk does not adopt filtering from the port configuration. You must still explicitly configure the filter on the port trunk. If you use the **show filter < *index* >** command for a filter created before the related source port was added to a trunk, the port number appears between asterisks (*), indicating that the filter action has been suspended for that filter. For example, if you create a filter on port 5, then create a trunk with ports 5 and 6, and display the results, you would see the following:

```
HP Switch(config)# filter source-port 5 drop 2
HP Switch(config)# trunk 5-6 trk1
HP Switch(config)# show filter

 Traffic/Security Filters

  IDX Filter Type  | Value
  --- ------------ + ------------
  1   Source Port  | *5*

HP Switch(config)# show filter 1

 Traffic/Security Filters

  Filter Type : Source Port
  Source Port : *5*

  Dest Port Type       | Action
  --------- --------- + -------
  1         100/1000T | Forward
  2         100/1000T | Forward
  3         100/1000T | Forward
  4         100/1000T | Forward
  .         .           .
  .         .           .
  .         .           .
```

The *5* shows that port 5 is configured for filtering, but the filtering action has been suspended while the port is a member of a trunk.

If you want the trunk to which port 5 belongs to filter traffic, then you must explicitly configure filtering on the trunk.

Note: If you configure an existing trunk for filtering and later add another port to the trunk, the switch will apply the filter to all traffic moving on any link in the trunk. If you remove a port from the trunk it returns to the configuration it had before it was added to the trunk

Figure 9-13. Example of Switch Response to Adding a Filtered Source Port to a Trunk

Editing a Source-Port Filter

The switch includes in one filter the action(s) for all destination ports and/or trunks configured for a given source port or trunk. Thus, if a source-port filter already exists and you want to change the currently configured action for some destination ports or trunks, use the **filter source-port** command to update the existing filter. For example, suppose you configure a filter to drop traffic received on port 8 and destined for ports 1 and 2. The resulting filter is shown on the left in figure 9-14. Later, you update the filter to drop traffic received on port 8 and destined for ports 3 through 5. Since only one filter exists for a given source port, the filter on traffic from port 8 appears as shown on the right in figure 9-14:

```
Switch(config)# show filter 1          Switch(config)# show filter 1
 Traffic/Security Filters               Traffic/Security Filters
  Filter Type : Source Port              Filter Type : Source Port
  Source Port : 8                        Source Port : 8

  Dest Port Type      | Action          Dest Port Type      | Action
  --------- --------- + -------         --------- --------- + -------
  1         100/1000T | Drop            1         100/1000T | Drop
  2         100/1000T | Drop            2         100/1000T | Drop
  3         100/1000T | Forward         3         100/1000T | Drop
  4         100/1000T | Forward         4         100/1000T | Drop
  5         100/1000T | Forward         5         100/1000T | Drop
  6         100/1000T | Forward         6         100/1000T | Forward
  7         100/1000T | Forward         7         100/1000T | Forward
```

Figure 9-14. Assigning Additional Destination Ports to an Existing Filter

```
HP Switch(config)# filter source-port 1 drop e d1-d4
HP Switch(config)# filter source-port 2 drop d1-d4
```

Figure 9-15. Configuring Various Traffic/Security Filters

Configuring a Multicast Traffic Filter

Syntax: [no] filter

[multicast < *mac- address* >]

*Specifies a multicast address. Inbound traffic received
(on any port) with this multicast address will be
filtered. (Default: Forward on all ports.)*

The **no** *form of the command deletes the multicast filter
for the < mac-address > multicast address and returns
the destination ports for that filter to the* **Forward** *action.*

[< forward | drop > < *port-list* >]

*Specifies whether the designated destination port(s)
should forward or drop the filtered traffic.*

For example, suppose you wanted to configure the filters in table 9-2 on a
switch.

Table 9-2.   Filter Example

Filter Type	Filter Value	Action	Destination Ports
Source-Port	Inbound ports: 1, 2*	Drop	D1-D4
Multicast	010000-123456	Drop	C1-C24, D5-D10
Multicast	010000-224466	Drop	B1-B4

*Because the switch allows one inbound port in a source-port filter, the requirement to filter ports 1 and 2 means you will configure two separate source-port filters.

The following commands configure the filters listed above:

```
HP Switch(config)# filter source-port 1 drop e d1-d4
HP Switch(config)# filter source-port 2 drop d1-d4
HP Switch(config)# filter multicast 010000-123456 drop e c1-c24,d5-d10
HP Switch(config)# filter multicast 010000-224466 drop e b1-b4
```

Figure 9-16. Configuring Various Traffic/Security Filters

Filter Indexing

The switch automatically assigns each new filter to the lowest-available index (IDX) number. The index numbers are included in the **show filter** command described in the next section and are used with the **show filter < *index* >** command to display detailed information about a specific filter.

If there are no filters currently configured, and you create three filters in succession, they will have index numbers 1 - 3. However, if you then delete the filter using index number "2" and then configure two new filters, the first new filter will receive the index number "2" and the second new filter will receive the index number "4". This is because the index number "2" was made vacant by the earlier deletion, and was therefore the lowest index number available for the next new filter.

Displaying Traffic/Security Filters

This command displays a listing of all filters by index number and also enables you to use the index number to display the details of individual filters.

Syntax: show filter

> *Lists the filters configured in the switch, with corresponding filter index (**IDX**) numbers.*
> **IDX**: *An automatically assigned index number used to identify the filter for a detailed information listing. A filter retains its assigned IDX number for as long as the filter exists in the switch. The switch assigns the lowest available IDX number to a new filter. This can result in a newer filter having a lower IDX number than an older filter if a previous filter deletion created a gap in the filter listing.*
> **Filter Type**: *Indicates the type of filter assigned to the **IDX** number (source-port, multicast, or protocol).*
> **Value**: *Indicates the port number or port-trunk name of the source port or trunk assigned to the filter*

[*index*]

> Lists the filter type and other data for the filter corresponding to the index number in the **show filter** output. Also lists, for each outbound destination port in the switch, the port number, port type, and filter action (forward or drop). The switch assigns the lowest available index number to a new filter. If you delete a filter, the index number for that filter becomes available for the next filter you create.

```
              HP Switch(config)# show filter              Lists all filters configured
                                                          in the switch.
                Traffic/Security Filters

                IDX  Filter Type  | Value
Filter Index Numbers   ---  -----------  +  ------------
(Automatically Assigned)    1   Source Port   1                Criteria for Individual
                1   Source Port   1                Filters
                2   Source Port   2
                3   Multicast     010000-123456
                4   Multicast     010000-224466


              HP Switch(config)# show filter              Uses the index number
                                                          (IDX) for a specific filter
                Traffic/Security Filters                  to list the details for that
                                                          filter only.
                Filter Type : Multicast
                Multicast Address : 010000-224466

                Dest Port  Type      | Action
                ---------  --------- + -------
                A1         1000LX    | Forward
                A2                   | Forward
                A3                   | Forward
                A4         1000SX    | Forward
                B1         100/1000T | Drop
                B2         100/1000T | Drop
                B3         100/1000T | Drop
                B4         100/1000T | Drop
                C1         10/100TX  | Forward
                C2         10/100TX  | Forward
                C3         10/100TX  | Forward
                 .          .           .
```

Figure 9-17. Example of Displaying Filter Data

10

Configuring Port-Based and Client-Based Access Control (802.1X)

Overview

Feature	Default	Menu	CLI	Web
Configuring Switch Ports as 802.1X Authenticators	Disabled	n/a	page 10-14	n/a
Configuring 802.1X Open VLAN Mode	Disabled	n/a	page 10-24	n/a
Configuring Switch Ports to Operate as 802.1X Supplicants	Disabled	n/a	page 10-38	n/a
Displaying 802.1X Configuration, Statistics, and Counters	n/a	n/a	page 10-41	n/a
How 802.1X Affects VLAN Operation	n/a	n/a	page 10-52	n/a
RADIUS Authentication and Accounting	Refer to "RADIUS Authentication, Authorization and Accounting" on page 5-1		

Why Use Port-Based or Client-Based Access Control?

Local Area Networks are often deployed in a way that allows unauthorized clients to attach to network devices, or allows unauthorized users to get access to unattended clients on a network. Also, the use of DHCP services and zero configuration make access to networking services easily available. This exposes the network to unauthorized use and malicious attacks. While access to the network should be made easy, uncontrolled and unauthorized access is usually not desirable. 802.1X simplifies security management by providing access control along with the ability to control user profiles from up to three RADIUS servers while allowing a given user to use the same entering valid user credentials for access from multiple points within the network.

General Features

802.1X on the HP switches covered in this manual 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 (Extensible Authentication Protocol) or CHAP (Challenge Handshake Authentication Protocol).

- Provision for enabling clients that do not have 802.1 supplicant software to use the switch as a path for downloading the software and initiating the authentication process (802.1X Open VLAN mode).

Client-Based access control option with support for up to 32 authenticated clients per-port.

Port-Based access control option allowing authentication by a single client to open the port. This option does not force a client limit and, on a port opened by an authenticated client, allows unlimited client access without requiring further authentication.

- Supplicant implementation using CHAP authentication and independent username and password configuration on each port.

■ Prevention of traffic flow in either direction on unauthorized ports.

■ Local authentication of 802.1X clients using the switch's local username and password (as an alternative to RADIUS authentication).

■ Temporary on-demand change of a port's VLAN membership status to support a current client's session. (This does not include ports that are members of a trunk.)

■ Session accounting with a RADIUS server, including the accounting update interval.

■ Use of Show commands to display session counters.

User Authentication Methods

The switch offers two methods for using 802.1X access control. Generally, the "Port Based" method supports one 802.1X-authenticated client on a port, which opens the port to an unlimited number of clients. The "Client-Based" method supports up to two 802.1X-authenticated clients on a port. In both cases, there are operating details to be aware of that can influence your choice of methods.

802.1X Client-Based Access Control

802.1X operation with access control on a per-client basis provides client-level security that allows LAN access to individual 802.1X clients (up to 32 per port), where each client gains access to the LAN by entering valid user credentials. This operation improves security by opening a given port only to individually authenticated clients, while simultaneously blocking access to the same port for clients that cannot be authenticated. All sessions must use the same

untagged VLAN. Also, an authenticated client can use any tagged VLAN memberships statically configured on the port, provided the client is configured to use the tagged VLAN memberships available on the port.

802.1X Port-Based Access Control

802.1X port-based access control provides port-level security that allows LAN access only on ports where a single 802.1X-capable client (supplicant) has entered authorized RADIUS user credentials. For reasons outlined below, this option is recommended for applications where only one client at a time can connect to the port. Using this option, the port processes all traffic as if it comes from the same client. Thus, in a topology where multiple clients can connect to the same port at the same time:

■   If the first client authenticates and opens the port, and then another client authenticates, the port responds as if the original client has initiated a reauthentication. With multiple clients authenticating on the port, the RADIUS configuration response to the latest client authentication replaces any other configuration from an earlier client authentication. If all clients use the same configuration this should not be a problem. But if the RADIUS server responds with different configurations for different clients, then the last client authenticated will effectively lock out any previously authenticated client. When *any* client to authenticate closes its session, the port will also close and remain so until another client successfully authenticates.

■   The most recent client authentication determines the untagged VLAN membership for the port. Also, any client able to use the port can access any tagged VLAN memberships statically configured on the port, provided the client is configured to use the available, tagged VLAN memberships.

■   If the first client authenticates and opens the port, and then one or more other clients connect without trying to authenticate, then the port configuration as determined by the original RADIUS response remains unchanged and all such clients will have the same access as the authenticated client. When the authenticated client closes the session, the port will also be closed to any other, unauthenticated clients that may have also been using the port.

This operation unblocks the port while an authenticated client session is in progress. In topologies where simultaneous, multiple client access is possible this can allow unauthorized and unauthenticated access by another client while an authenticated client is using the port. If you want to allow only authenticated clients on the port, then client-based access control (page 10-2) should be used instead of port-based access control. Using the client-based method enables you to specify up to 32 authenticated clients.

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

Providing a Path for Downloading 802.1X Supplicant Software. For clients that do not have the necessary 802.1X supplicant software, there is also the option to configure the 802.1X Open VLAN mode. This mode allows you to assign such clients to an isolated VLAN through which you can provide the necessary supplicant software these clients need to begin the authentication process. (Refer to "802.1X Open VLAN Mode" on page 10-24.)

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.
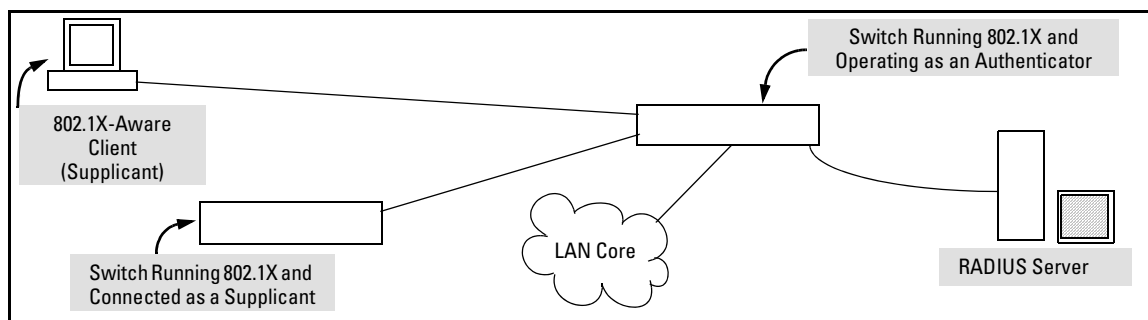
Figure 10-1. Example of an 802.1X Application

Accounting . The switch also provides RADIUS Network accounting for 802.1X access. Refer to "RADIUS Authentication, Authorization and Accounting" on page 5-1.

Terminology

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

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

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 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 applications, a switch that requires a supplicant to provide the proper credentials before being allowed access to the network.

CHAP (MD5): Challenge Handshake Authentication Protocol.

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

Client-Based Authentication: The 802.1X extension in the switches covered in this guide. In this operation, multiple clients on the same port must individually authenticate themselves.

Guest VLAN: See "Unauthorized-Client VLAN".

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

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

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

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

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

Port-Based Authentication: In this operation, the first client on a port to authenticate itself unblocks the port for the duration of the client's 802.1X-authenticated session. The switches covered in this guide use port-based authentication.

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

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

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

Unauthorized-Client VLAN: A conventional, static VLAN statically configured on the switch. It is used to provide access to a client prior to authentication, and is sometimes termed a *guest* VLAN. It should be set up to allow an unauthenticated client to access only the initialization services necessary to establish an authenticated connection, plus any other desirable services whose use by an unauthenticated client poses no security threat to your network. (Note that an unauthenticated client has access to all network resources that have membership in the VLAN you

designate as the Unauthorized-Client VLAN.) A port configured to use a given Unauthorized-Client VLAN does not have to be statically configured as a member of that VLAN as long as at least one other port on the switch is statically configured as a tagged or untagged member of the same Unauthorized-Client VLAN. An unauthorized-client VLAN is available on a port only if there is no authenticated client already using the port.

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

General 802.1X Authenticator Operation

This operation provides security on a direct, point-to-point link between a single client and the switch, where both devices are 802.1X-aware. (If you expect desirable clients that do not have the necessary 802.1X supplicant software, you can provide a path for downloading such software by using the 802.1X Open VLAN mode—refer to "802.1X Open VLAN Mode" on page 10-24.)

Example of the Authentication Process

Suppose that you have configured a port on the switch for 802.1X authentication operation. If you then connect an 802.1X-aware client (supplicant) to the port and attempt to log on:

1. When the switch detects the client on the port, it blocks access to the LAN from that port.

2. The switch responds with an identity request.

3. The client responds with a user name that uniquely defines this request for the client.

4. The switch responds in one of the following ways:

- If 802.1X (port-access) on the switch is configured for RADIUS authentication, the switch then forwards the request to a RADIUS server.
    i.   The server responds with an access challenge which the switch forwards to the client.
    ii. The client then provides identifying credentials (such as a user certificate), which the switch forwards to the RADIUS server.
    iii. The RADIUS server then checks the credentials provided by the client.
    iv. If the client is successfully authenticated and authorized to connect to the network, then the server notifies the switch to allow access to the client. Otherwise, access is denied and the port remains blocked.
- If 802.1X (port-access) on the switch is configured for local authentication, then:
    i. The switch compares the client's credentials with the username and password configured in the switch (Operator or Manager level).
    ii. If the client is successfully authenticated and authorized to connect to the network, then the switch allows access to the client. Otherwise, access is denied and the port remains blocked.

Switch-Port Supplicant Operation

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

- Switch "A" has port 1 configured for 802.1X supplicant operation.

- You want to connect port 1 on switch "A" to port 5 on switch "B".
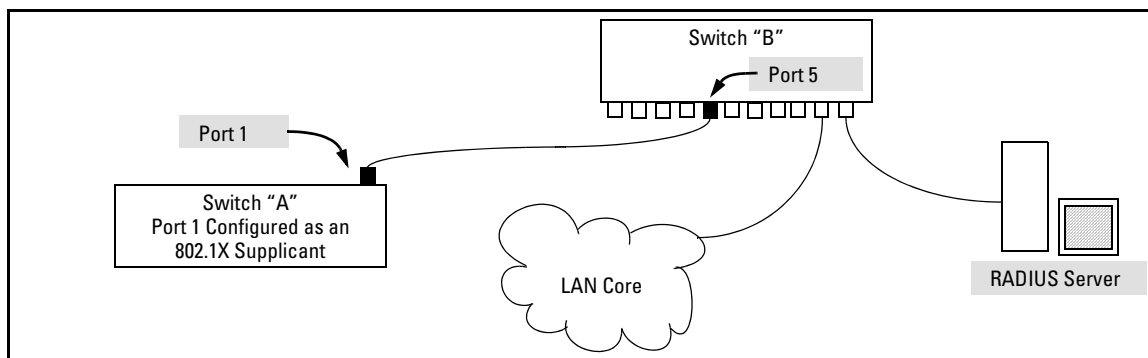
Figure 10-2. Example of Supplicant Operation

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

   - If, after the supplicant port sends the configured number of start packets, it does not receive a response, it assumes that switch "B" is not 802.1X-aware, and transitions to the authenticated state. If switch "B" is operating properly and is not 802.1X-aware, then the link should begin functioning normally, but without 802.1X security.

   - If, after sending one or more start packets, port 1 receives a request packet from port 5, then switch "B" is operating as an 802.1X authenticator. The supplicant port then sends a response/ID packet. Switch "B" forwards this request to a RADIUS server.

2. The RADIUS server then responds with an MD5 access challenge that switch "B" forwards to port 1 on switch "A".

3. Port 1 replies with an MD5 hash response based on its username and password or other unique credentials. Switch "B" forwards this response to the RADIUS server.

4. The RADIUS server then analyzes the response and sends either a "success" or "failure" packet back through switch "B" to port 1.

   - A "success" response unblocks port 5 to normal traffic from port 1.

   - A "failure" response continues the block on port 5 and causes port 1 to wait for the "held-time" period before trying again to achieve authentication through port 5.

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

General Operating Rules and Notes

■   In the client-based mode, when there is an authenticated client on a port, the following traffic movement is allowed:

   - Multicast and broadcast traffic is allowed on the port.

   - Unicast traffic to authenticated clients on the port is allowed.

   - All traffic from authenticated clients on the port is allowed.

■ 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.

■ Using client-based 802.1X authentication, when a port on the switch is configured as an authenticator the port allows only authenticated clients up to the currently configured client limit.

For clients that do not have the proper 802.1X supplicant software, the optional 802.1X Open VLAN mode can be used to open a path for downloading 802.1X supplicant software to a client or to provide other services for unauthenticated clients.

■ Using port-based 802.1X authentication, when a port on the switch is configured as an authenticator, one authenticated client opens the port. Other clients that are not running an 802.1X supplicant application can have access to the switch and network through the opened port. If another client uses an 802.1X supplicant application to access the opened port, then a re-authentication occurs using the RADIUS configuration response for the latest client to authenticate. To control access by all clients, use the client-based method.

■ If a port on switch "A" is configured as an 802.1X supplicant and is connected to a port on another switch, "B", that is not 802.1X-aware, access to switch "B" will occur without 802.1X security protection.

■ You can configure a port as both an 802.1X authenticator *and* an 802.1X supplicant.

■ If a port on switch "A" is configured as both an 802.1X authenticator *and* supplicant and is connected to a port on another switch, "B", that is not 802.1X-aware, access to switch "B" will occur without 802.1X security protection, but switch "B" will not be allowed access to switch "A".

■ If a client already has access to a switch port when you configure the port for 802.1X authenticator operation, the port will block the client from further network access until it can be authenticated.

■ On a port configured for 802.1X with RADIUS authentication, if the RADIUS server specifies a VLAN for the supplicant and the port is a trunk member, the port will be blocked. If the port is later removed from the trunk, the port will try to authenticate the supplicant. If authentication is successful, the port becomes unblocked. Similarly, if the supplicant is authenticated and later the port becomes a trunk member, the port will be blocked. If the port is then removed from the trunk, it tries to re-authenticate the supplicant. If successful, the port becomes unblocked.

■ To help maintain security, 802.1X and LACP cannot both be enabled on the same port. If you try to configure 802.1X on a port already configured for LACP (or the reverse) you will see a message similar to the following:

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

General Setup Procedure for 802.1X Access Control

Do These Steps Before You Configure 802.1X Operation

1. Configure a local username and password on the switch for both the Operator (login) and Manager (enable) access levels. (While this may or may not be required for your 802.1X configuration, HP recommends that you use a local username and password pair at least until your other security measures are in place.)

> ***Syntax:*** password <operator | manager | all> [user-name <*name*>] <*password* >
>
> > *Configures the operator username and password used to access the network through 802.1X authentication.*
> >
> > user-name <*name*>
> >
> > > *Operator username (text string) used only for local authentication of 802.1X clients. This value is different from the local operator user-name configured with the **password** command for management access.*
> >
> > <*password*>
> >
> > > *Operator password (text string) used only for local authentication of 802.1X clients. This value is different from the local operator pass-word configured with the **password** command for management access.*

2. Determine which ports on the switch you want to operate as authenticators and/or supplicants, and disable LACP on these ports. (See the "" on page 10-11.)

   To display the current configuration of 802.1X, Web-based, and MAC authentication on all switch ports, enter the **show port-access config** command.

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

 Port Access Status Summary
  Port-access authenticator activated [No] : Yes
  Allow RADIUS-assigned dynamic (GVRP) VLANs [No] : Yes


       Supplicant Authenticator Web Auth Mac Auth
  Port Enabled    Enabled       Enabled  Enabled
  ---- ---------- ------------- -------- --------
  1    Yes        No            No       No
  2    No         No            No       Yes
  3    No         Yes           No       No
  4    No         No            No       No
  5    No         No            No       No
     ...
```

Figure 10-3. Example of show port-access config Command Output

3. Determine whether to use client-based access control or port-based access control.

4. Determine whether to use the optional 802.1X Open VLAN mode for clients that are not 802.1X-aware; that is, for clients that are not running 802.1X supplicant software. (This will require you to provide downloadable software that the client can use to enable an authentication session.) For more on this topic, refer to "802.1X Open VLAN Mode" on page 10-24.

5. 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.)

6. Unless you are using only the switch's local username and password for 802.1X authentication, configure at least one RADIUS server to authenticate access requests coming through the ports on the switch from external supplicants (including switch ports operating as 802.1X supplicants). You

can use up to three RADIUS servers for authentication; one primary and two backups. Refer to the documentation provided with your RADIUS application.

Overview: Configuring 802.1X Authentication on the Switch

This section outlines the steps for configuring 802.1X on the switch. For detailed information on each step, refer to "RADIUS Authentication, Authorization and Accounting" on page 5-1 or "Configuring Switch Ports To Operate As Supplicants for 802.1X Connections to Other Switches" on page 10-38.

1. Enable 802.1X authentication on the individual ports you want to serve as authenticators. On the ports you will use as authenticators, either accept the default 802.1X settings or change them, as necessary. Note that, by default, the port-control parameter is set to **auto** for all ports on the switch. This requires a client to support 802.1X authentication and to provide valid credentials to get network access. Refer to page 10-14.

2. If you want to provide a path for clients without 802.1X supplicant software to download the software so that they can initiate an authentication session, enable the 802.1X Open VLAN mode on the ports you want to support this feature. Refer to page 10-24.

3. Configure the 802.1X authentication type. Options include:
   - Local Operator username and password (the default). This option allows a client to use the switch's local username and password as valid 802.1X credentials for network access.
   - EAP RADIUS: This option requires your RADIUS server application to support EAP authentication for 802.1X.
   - CHAP (MD5) RADIUS: This option requires your RADIUS server application to support CHAP (MD5) authentication.

   See page 10-19.

4. If you select either **eap-radius** or **chap-radius** for step 3, use the **radius host** command to configure up to three RADIUS server IP address(es) on the switch. See page 10-20.

5. Enable 802.1X authentication on the switch. See page 10-14.

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

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

7. If you are using Port Security on the switch, configure the switch to allow only 802.1X access on ports configured for 802.1X operation, and (if desired) the action to take if an unauthorized device attempts access through an 802.1X port. See page 10-36.

8. If you want a port on the switch to operate as a supplicant in a connection with a port operating as an 802.1X authenticator on another device, then configure the supplicant operation. (Refer to "Configuring Switch Ports To Operate As Supplicants for 802.1X Connections to Other Switches" on page 10-38.)

Configuring Switch Ports as 802.1X Authenticators

802.1X Authentication Commands	Page
[no] aaa port-access authenticator < [ethernet] < *port-list* >	10-14
[control	quiet-period	tx-period	client-limit	supplicant-timeout	server-timeout	logoff-period	max-requests	reauth-period	auth-vid	unauth-vid	initialize	reauthenticate	clear-statistics]	10-15
aaa authentication port-access	10-19
< local	eap-radius	chap-radius >	
[no] aaa port-access authenticator active	10-14
[no] port-security [ethernet] < *port-list* > learn-mode port-access	10-36
802.1X Open VLAN Mode Commands	10-24
802.1X Supplicant Commands	10-38
802.1X-Related Show Commands	10-41
RADIUS server configuration	10-20

1. Enable 802.1X Authentication on Selected Ports

This task configures the individual ports you want to operate as 802.1X authenticators for point-to-point links to 802.1X-aware clients or switches, and consists of two steps:

A. Enable the selected ports as authenticators.

B. Specify either client-based or port-based 802.1X authentication.

(Actual 802.1X operation does not commence until you perform step 5 on page 10-13 to activate 802.1X authentication on the switch.)

Note

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

A. Enable the Selected Ports as Authenticators and Enable the (Default) Port-Based Authentication

Syntax: [no] aaa port-access authenticator < *port-list* >

> *Enables specified ports to operate as 802.1X authenticators and enables port-based authentication. (To enable client-based authentication, execute this command first, and then execute the client-limit < port-list > version of this command described in the next section.) The* **no** *form of the command removes 802.1X authentication from < port-list >. To activate configured 802.1X operation, you must enable 802.1X authentication. Refer to "5. Enable 802.1X Authentication on the switch" on page 10-13.*

B. Specify Client-Based or Return to Port-Based 802.1X Authentication

Client-Based 802.1X Authentication.

Syntax: aaa port-access authenticator client-limit < *port-list* > < 1 - 2>

> *Used after executing* **aaa port-access authenticator < *port-list* >** *(above) to convert authentication from port-based to client-based. Specifies client-based 802.1X authentication and the maximum number of 802.1X-authenticated client sessions allowed on each of the ports in* **< *port-list* >**. *If a port currently has no authenticated client sessions, the next authenticated client session the port accepts determines the untagged VLAN membership to which the port is assigned during the session. If another client session begins later on the same port while an earlier session is active, the later session will be on the same untagged VLAN membership as the earlier session.*

Port-Based 802.1X Authentication.

Syntax: no aaa port-access authenticator client-limit

Used to convert a port from client-based authentication to port-based authentication, which is the default setting for ports on which authentication is enabled. (Executing **aaa port-access authenticator < *port-list* >** *enables 802.1X authentication on* **< *port-list* >** *and enables port-based authentication—page 10-14.) If a port currently has no authenticated client sessions, the next authenticated client session the port accepts determines the untagged VLAN membership to which the port is assigned during the session. If another authenticated client session begins later on the same port while an earlier session is active, the later session replaces the currently active session and will be on the untagged VLAN membership specified by the RADIUS server for the later session.*

Example: Configuring Client-Based 802.1X Authentication

This example enables ports 10-12 to operate as authenticators, and then configures the ports for client-based authentication.

```
HP Switch(config)# aaa port-access authenticator 10-12
HP Switch(config)# aaa port-access authenticator 10-12 client-limit 2
```

Figure 10-4. Example of Configuring Client-Based 802.1X Authentication

Example: Configuring Port-Based 802.1X Authentication

This example enables ports 13-15 to operate as authenticators, and then configures the ports for port-based authentication.

```
HP Switch(config)# aaa port-access authenticator 13-15
```

Figure 10-5. Example of Configuring Port-Based 802.1X Authentication

2. Reconfigure Settings for Port-Access

The commands in this section are initially set by default and can be reconfigured as needed.

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

[control < authorized | auto | unauthorized >]

Controls authentication mode on the specified port:

authorized: *Also termed "<u>Force Authorized</u>". Gives access to a device connected to the port. In this case, the device does not have to provide 802.1X credentials or support 802.1X authentication. (You can still configure console, Telnet, or SSH security on the port.)*

auto *(the default): The device connected to the port must support 802.1X authentication and provide valid credentials to get network access. (Optional: You can use the Open VLAN mode to provide a path for clients without 802.1X supplicant software to down-load this software and begin the authentication process. Refer to "802.1X Open VLAN Mode" on page 10-24.)*

unauthorized: *Also termed "<u>Force Unauthorized</u>". 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 there is no response within the configured time frame, the switch assumes that the authentication attempt has timed out. Depending on the current **max-requests** *setting, the switch will either send a new request to the server or end the authentication session. (Default: 30 seconds)*

[max-requests < 1 - 10 >]

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

[reauth-period < 0 - 9999999 >]

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

[unauth-vid < *vlan-id* >]

Configures an existing static VLAN to be the Unauthorized-Client VLAN. This enables you to provide a path for clients without supplicant software to download the software and begin an authentication session. Refer to "802.1X Open VLAN Mode" on page 10-24.

[logoff-period]< 1 - 999999999 >

Configures the period of time the switch waits for client activity before removing an inactive client from the port. (Default: 300 seconds)

[auth-vid < *vid* >]

Configures an existing, static VLAN to be the Authorized-Client VLAN. Refer to "802.1X Open VLAN Mode" on page 10-24.

[initialize]

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

[reauthenticate]

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

[clear-statistics]

Clears authenticator statistics counters.

3. Configure the 802.1X Authentication Method

To specify how the switch will authenticate the credentials provided by a supplicant connected to a switch port configured as an 802.1X authenticator.

You can configure **local**, **chap-radius** or **eap-radius** as the primary password authentication method for the port-access method. You also need to select **none or authorized** as a secondary, or backup, method.

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

Configures **local**, **chap-radius (**MD5**),** *or* **eap-radius** *as the primary password authentication method for port-access. The default primary authentication is* **local***. (Refer to the documentation for your RADIUS server application.)*

For switches covered in this guide, you must use the **password port-access** *command to configure the operator username and password for 802.1X access. See "General Setup Procedure for 802.1X Access Control" on page 10-11 for more information.*

[<none | authorized >]

Provides options for secondary authentication. The **none** *option specifies that a backup authentication method is not used. The* **authorized** *option allows access without authentication. (default:* **none***).*

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

```
HP Switch(config)# aaa authentication port-access eap-radius
HP Switch(config)# show auth
  Status and Counters - Authentication Information
    Login Attempts : 3
    Respect Privilege : Disabled


                | Login      Login     Enable     Enable
    Access Task | Primary    Secondary Primary    Secondary
    ----------- + --------- ---------- ---------- ---------
    Console     | Local      None      Local      None
    Telnet      | Local      None      Local      None
    Port-Access | EapRadius
    Webui       | Local      None      Local      None
    SSH         | Local      None      Local      None
    Web-Auth    | ChapRadius
    MAC-Auth    | ChapRadius
```

Figure 10-6. Example of 802.1X (Port-Access) Authentication

4. Enter the RADIUS Host IP Address(es)

If you selected either **eap-radius** or **chap-radius** for the authentication method, configure the switch to use 1 to 3 RADIUS servers for authentication. The following syntax shows the basic commands. For coverage of all commands related to RADIUS server configuration, refer to "RADIUS Authentication, Authorization and Accounting" on page 5-1.

Syntax: radius host < *ip-address* >

Adds a server to the RADIUS configuration.

[key < *server-specific key-string* >]

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

radius-server key < *global key-string* >

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

5. Enable 802.1X Authentication on the Switch

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

Syntax: aaa port-access authenticator active

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

6. Optionally Resetting Authenticator Operation

After authentication has begun operating, these commands can be used to reset authentication and related statistics on specific ports.

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

[initialize]

*On the specified ports, blocks inbound and outbound traffic and restarts the 802.1X authentication process. This happens only on ports configured with **control auto** and actively operating as 802.1X authenticators.*

[reauthenticate]

On the specified ports, forces reauthentication (unless the authenticator is in "HELD" state).

[clear-statistics]

On the specified ports, clears authenticator statistics counters.

7. Optional: Configure 802.1X Controlled Directions

After you enable 802.1X authentication on specified ports, you can use the **aaa port-access controlled-directions** command to configure how a port transmits traffic before it successfully authenticates a client and enters the authenticated state.

As documented in the IEEE 802.1X standard, an 802.1X-aware port that is unauthenticated can control traffic in either of the following ways:

■   In both ingress and egress directions by disabling both the reception of incoming frames and transmission of outgoing frames

■   Only in the ingress direction by disabling only the reception of incoming frames.

Prerequisite. As documented in the IEEE 802.1X standard, the disabling of incoming traffic and transmission of outgoing traffic on an 802.1X-aware egress port in an unauthenticated state (using the **aaa port-access controlled-directions in** command) is supported only if:

■ The port is configured as an edge port in the network using the **spanning-tree edge-port** command.

■ The 802.1s Multiple Spanning Tree Protocol (MSTP) or 802.1w Rapid Spanning Tree Protocol (RSTP) is enabled on the switch. MSTP and RSTP improve resource utilization while maintaining a loop-free network.

For information on how to configure the prerequisites for using the **aaa port-access controlled-directions in** command, see Chapter 4, "Multiple Instance Spanning-Tree Operation" in the *Advanced Traffic Management Guide*.

Syntax: aaa port-access <*port-list*> controlled-directions <both | in>

> **both** *(default)***:** *Incoming and outgoing traffic is blocked on an 802.1X-aware port before authentication occurs.*
> **in**: *Incoming traffic is blocked on an 802.1X-aware port before authentication occurs. Outgoing traffic with unknown destination addresses is flooded on unauthenticated 802.1X-aware ports.*

Wake-on-LAN Traffic

The Wake-on-LAN feature is used by network administrators to remotely power on a sleeping workstation (for example, during early morning hours to perform routine maintenance operations, such as patch management and software updates).

The **aaa port-access controlled-direction in** command allows Wake-on-LAN traffic to be transmitted on an 802.1X-aware egress port that has not yet transitioned to the 802.1X authenticated state; the **controlled-direction both** setting prevents Wake-on-LAN traffic to be transmitted on an 802.1X-aware egress port until authentication occurs.

N o t e   Although the **controlled-direction in** setting allows Wake-on-LAN traffic to traverse the switch through unauthenticated 802.1X-aware egress ports, it does not guarantee that the Wake-on-LAN packets will arrive at their destination. For example, firewall rules on other network devices and VLAN rules may prevent these packets from traversing the network.

Operating Notes

■ Using the **aaa port-access controlled-directions in** command, you can enable the transmission of Wake-on-LAN traffic on unauthenticated egress ports that are configured for any of the following port-based security features:

- 802.1X authentication
- MAC authentication
- Web authentication

Because a port can be configured for more than one type of authentication to protect the switch from unauthorized access, the last setting you configure with the **aaa port-access controlled-directions** command is applied to all authentication methods configured on the switch.

For information about how to configure and use MAC and Web authentication, refer to chapter 3, "Web and MAC Authentication".

■ To display the currently configured 802.1X Controlled Directions value, enter the **show port-access authenticator config** command as shown in Figure 10-10.

■ When an 802.1X-authenticated port is configured with the **controlled-directions in** setting, eavesdrop prevention is not supported on the port.

Example: Configuring 802.1X Controlled Directions

The following example shows how to enable the transmission of Wake-on-LAN traffic in the egress direction on an 802.1X-aware port before it transitions to the 802.1X authenticated state and successfully authenticates a client device.

```
HP Switch(config)# aaa port-access authenticator 10
HP Switch(config)# aaa authentication port-access eap-radius
HP Switch(config)# aaa port-access authenticator active
HP Switch(config)# aaa port-access 10 controlled-directions in
```

Figure 10-7. Example of Configuring 802.1X Controlled Directions

802.1X Open VLAN Mode

Introduction

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

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

■ Acquiring IP addressing from a DHCP server

■ Downloading the 802.1X supplicant software necessary for an authentication session

The 802.1X Open VLAN mode solves this problem by temporarily suspending the port's static, tagged and untagged VLAN memberships and placing the port in a designated *Unauthorized-Client VLAN*. In this state the client can proceed with initialization services, such as acquiring IP addressing and 802.1X software, and starting the authentication process.

Note

On ports configured to allow multiple sessions using 802.1X client-based access control, all clients must use the same untagged VLAN. On a given port where there are no currently active, authenticated clients, the first *authenticated* client determines the untagged VLAN in which the port will operate for all subsequent, overlapping client sessions.

If the switch operates in an environment where some valid clients will not be running 802.1X supplicant software and need to download it from your network. Then, because such clients would need to use the Unauthorized-Client VLAN and authenticated clients would be using a different VLAN (for security reasons), allowing multiple clients on an 802.1X port can result in blocking some or all clients needing to use the Unauthorized-Client VLAN.

On ports configured for port-based 802.1X access control, if multiple clients try to authenticate on the same port, the most recently authenticated client determines the untagged VLAN membership for that port. Clients that connect without trying to authenticate will have access to the untagged VLAN membership that is currently assigned to the port.

VLAN Membership Priorities

Following client authentication, an 802.1X port resumes membership in any tagged VLANs for which it is already assigned in the switch configuration. The port also becomes an untagged member of one VLAN according to the following order of options:

1. **1st Priority:** The port joins a VLAN to which it has been assigned by a RADIUS server during authentication.

2. **2nd Priority:** If RADIUS authentication does not include assigning the port to a VLAN, then the switch assigns the port to the VLAN entered in the port's 802.1X configuration as an *Authorized-Client* VLAN, if configured.

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

A port assigned to a VLAN by an Authorized-Client VLAN configuration (or a RADIUS server) will be an untagged member of the VLAN for the duration of the authenticated session. This applies even if the port is also configured in the switch as a tagged member of the same VLAN.

N o t e

After client authentication, the port resumes membership in any tagged
VLANs for which it is configured. If the port belongs to a tagged VLAN used
for 1 or 2 above, then it operates as an untagged member of that VLAN while
the client is connected. When the client disconnects, the port reverts to tagged
membership in the VLAN.

Use Models for 802.1X Open VLAN Modes

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

■ **Unauthorized-Client VLAN:** Configure this VLAN when unauthen-
ticated, friendly clients will need access to some services before being
authenticated.

■ **Authorized-Client VLAN:** Configure this VLAN for authenticated
clients when the port is not statically configured as an untagged
member of a VLAN you want clients to use, or when the port is
statically configured as an untagged member of a VLAN you do not
want clients to use. (A port can be configured as untagged on only
one VLAN. When an Authorized-Client VLAN is configured, it will
always be untagged and will block the port from using a statically
configured, untagged membership in another VLAN.) Note that after
client authentication, the port returns to membership in any tagged
VLANs for which you have configured it. See the "Note", above.

Table 10-1.   802.1X Open VLAN Mode Options

802.1X Per-Port Configuration	Port Response
No Open VLAN mode:	The port automatically blocks a client that cannot initiate an authentication session.
Open VLAN mode with **both** of the following configured:	
Unauthorized-Client VLAN	• When the port detects a client, it automatically becomes an untagged member of this VLAN. If you previously configured the port as a static, tagged member of the VLAN, membership temporarily changes to untagged while the client remains unauthenticated.
	• If the port already has a statically configured, untagged membership in another VLAN, then the port temporarily closes access to this other VLAN while in the Unauthorized-Client VLAN.
	• To limit security risks, the network services and access available on the Unauthorized-Client VLAN should include only what a client needs to enable an authentication session. If the port is statically configured as a tagged member of any other VLANs, access to these VLANs is blocked while the port is a member of the Unauthorized-Client VLAN.

Table 10-1.   802.1X Open VLAN Mode Options

802.1X Per-Port Configuration	Port Response
Authorized-Client VLAN	• After the client is authenticated, the port drops membership in the Unauthorized-Client VLAN and becomes an untagged member of this VLAN. **Note:** if RADIUS authentication assigns a VLAN, the port temporarily becomes a member of the RADIUS-assigned VLAN —instead of the Authorized-Client VLAN—while the client is connected. • If the port is statically configured as a tagged member of a VLAN, and this VLAN is used as the Authorized-Client VLAN, then the port temporarily becomes an untagged member of this VLAN when the client becomes authenticated. When the client disconnects, the port returns to tagged membership in this VLAN. • If the port is statically configured as a tagged member of a VLAN that is not used by 802.1X Open VLAN mode, the port returns to tagged membership in this VLAN upon successful authentication. This happens even if the RADIUS server assigns the port to another, authorized VLAN. If the port is already configured as a tagged member of a VLAN that RADIUS assigns as an authorized VLAN, then the port becomes an untagged member of that VLAN for the duration of the client connection. After the client disconnects, the port returns to tagged membership in that VLAN.
Open VLAN Mode with **Only** an **Unauthorized-Client VLAN** Configured:	• When the port detects a client, it automatically becomes an untagged member of this VLAN. To limit security risks, the network services and access available on this VLAN should include only what a client needs to enable an authentication session. If the port is statically configured as an untagged member of another VLAN, the switch temporarily removes the port from membership in this other VLAN while membership in the Unauthorized-Client VLAN exists. • After the client is authenticated, and if the port is statically configured as an untagged member of another VLAN, the port's access to this other VLAN is restored. **Note:** If RADIUS authentication assigns a VLAN to the port, this assignment overrides any statically configured, untagged VLAN membership on the port (while the client is connected). • If the port is statically configured as a tagged member of a VLAN that is not used by 802.1X Open VLAN mode, the port returns to tagged membership in this VLAN upon successful client authentication. This happens even if the RADIUS server assigns the port to another, authorized VLAN. Note that if the port is already configured as a tagged member of a VLAN that RADIUS assigns as an authorized VLAN, then the port becomes an untagged member of that VLAN for the duration of the client connection. After the client disconnects, the port returns to tagged membership in that VLAN.
Open VLAN Mode with **Only** an **Authorized-Client VLAN** Configured:	

Table 10-1.   802.1X Open VLAN Mode Options

802.1X Per-Port Configuration	Port Response
	• Port automatically blocks a client that cannot initiate an authentication session.
	• If the client successfully completes an authentication session, the port becomes an untagged member of this VLAN. **Note:** if RADIUS authentication assigns a VLAN, the port temporarily becomes an untagged member of the RADIUS-assigned VLAN —instead of the Authorized-Client VLAN—while the client is connected.
	• If the port is statically configured as a tagged member of any other VLAN, the port returns to tagged membership in this VLAN upon successful client authentication. This happens even if the RADIUS server assigns the port to another, authorized VLAN. If the port is already configured as a tagged member of a VLAN that RADIUS assigns as an authorized VLAN, then the port becomes an untagged member of that VLAN for the duration of the client connection. After the client disconnects, the port returns to tagged membership in that VLAN.

Operating Rules for Authorized-Client and Unauthorized-Client VLANs

Condition	Rule
Static VLANs used as *Authorized-Client* or *Unauthorized-Client* VLANs	These must be configured on the switch before you configure an 802.1X authenticator port to use them. (Use the **vlan** < ***vlan-id*** > command or the VLAN Menu screen in the Menu interface.)
VLAN Assignment Received from a RADIUS Server	If the RADIUS server specifies a VLAN for an authenticated supplicant connected to an 802.1X authenticator port, this VLAN assignment overrides any Authorized-Client VLAN assignment configured on the authenticator port. This is because both VLANs are untagged, and the switch allows only one untagged VLAN membership per-port. For example, suppose you configured port 4 to place authenticated supplicants in VLAN 20. If a RADIUS server authenticates supplicant "A" and assigns this supplicant to VLAN 50, then the port can access VLAN 50 as an untagged member while the client session is running. When the client disconnects from the port, then the port drops these assignments and uses the untagged VLAN memberships for which it is statically configured. (After client authentication, the port resumes any tagged VLAN memberships for which it is already configured. For details, refer to the Note on page 10-26.)
Temporary VLAN Membership During a Client Session	• Port membership in a VLAN assigned to operate as the Unauthorized-Client VLAN is temporary, and ends when the client receives authentication or the client disconnects from the port, whichever is first. • Port membership in a VLAN assigned to operate as the Authorized-Client VLAN is also temporary, and ends when the client disconnects from the port. If a VLAN assignment from a RADIUS server is used instead, the same rule applies.

Condition	Rule
Effect of Unauthorized-Client VLAN session on untagged port VLAN membership	• When an unauthenticated client connects to a port that is already configured with a static, untagged VLAN, the switch temporarily moves the port to the Unauthorized-Client VLAN (also untagged). (While the Unauthorized-Client VLAN is in use, the port does not access the static, untagged VLAN.) • When the client either becomes authenticated or disconnects, the port leaves the Unauthorized-Client VLAN and reacquires its untagged membership in the statically configured VLAN.
Effect of Authorized-Client VLAN session on untagged port VLAN membership.	• When a client becomes authenticated on a port that is already configured with a static, untagged VLAN, the switch temporarily moves the port to the Authorized-Client VLAN (also untagged). While the Authorized-Client VLAN is in use, the port does not have access to the statically configured, untagged VLAN. • When the authenticated client disconnects, the switch removes the port from the Authorized-Client VLAN and moves it back to the untagged membership in the statically configured VLAN. (After client authentication, the port resumes any tagged VLAN memberships for which it is already configured. For details, refer to the Note on page 10-26.)
Multiple Authenticator Ports Using the Same Unauthorized-Client and Authorized-Client VLANs	You can use the same static VLAN as the Unauthorized-Client VLAN for all 802.1X authenticator ports configured on the switch. Similarly, you can use the same static VLAN as the Authorized-Client VLAN for all 802.1X authenticator ports configured on the switch. *Caution: Do not use the same static VLAN for both the unauthorized and the Authorized-Client VLAN. Using one VLAN for both creates a security risk by defeating the isolation of unauthenticated clients.*
Effect of Failed Client Authentication Attempt	When there is an Unauthorized-Client VLAN configured on an 802.1X authenticator port, an unauthorized client connected to the port has access only to the network resources belonging to the Unauthorized-Client VLAN. This access continues until the client disconnects from the port. (If there is no Unauthorized-Client VLAN configured on the authenticator port, the port simply blocks access for any unauthorized client that cannot be authenticated.)
IP Addressing for a Client Connected to a Port Configured for 802.x Open VLAN Mode	A client can either acquire an IP address from a DHCP server or have a preconfigured, manual IP address before connecting to the switch.
802.1X Supplicant Software for a Client Connected to a Port Configured for 802.1X Open VLAN Mode	A friendly client, without 802.1X supplicant software, connecting to an authenticator port must be able to download this software from the Unauthorized-Client VLAN before authentication can begin.

Condition	Rule

Switch with a Port Configured To Allow Multiple Authorized-Client Sessions	When a new client is authenticated on a given port:
	• If no other clients are authenticated on that port, then the port joins one VLAN in the following order of precedence:
	   a. A RADIUS-assigned VLAN, if configured.
	   b. An Authenticated-Client VLAN, if configured.
	   c. A static, port-based VLAN to which the port belongs as an untagged member.
	   d. Any VLAN(s) to which the port is configured as a tagged member (provided that the client can operate in that VLAN).
	• If another client is already authenticated on the port, then the port is already assigned to a VLAN for the previously-existing client session, and the new client must operate in this same VLAN, regardless of other factors. (This means that a client without 802.1X client authentication software cannot access a configured, Unauthenticated-Client VLAN if another, authenticated client is already using the port.)

Note: Limitation on Using an Unauthorized-Client VLAN on an 802.1X Port Configured to Allow Multiple-Client Access	You can optionally enable switches to allow up to 32 clients per-port. The Unauthorized-Client VLAN feature can operate on an 802.1X-configured port regardless of how many clients the port is configured to support. However, all clients on the same port must operate through the same untagged VLAN membership. This means that any client accessing a given port must be able to authenticate and operate on the same VLAN as any other previously authenticated clients that are currently using the port. Thus, an Unauthorized-Client VLAN configured on a switch port that allows multiple 802.1X clients cannot be used if there is already an authenticated client using the port on another VLAN. Also, a client using the Unauthenticated-Client VLAN will be blocked when another client becomes authenticated on the port. For this reason, the best utilization of the Unauthorized-Client VLAN feature is in instances where only one client is allowed per-port. Otherwise, unauthenticated clients are subject to being blocked at any time by authenticated clients using a different VLAN. (Using the same VLAN for authenticated and unauthenticated clients can create a security risk and is not recommended.)

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

Setting Up and Configuring 802.1X Open VLAN Mode

Preparation. This section assumes use of both the Unauthorized-Client and Authorized-Client VLANs. Refer to Table 10-3 on page 10-49 for other options.

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

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

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

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

Note that if an 802.1X authenticator port is an untagged member of another VLAN, the port's access to that other VLAN will be temporarily removed while an authenticated client is connected to the port. For example, if:
   i.   Port 5 is an untagged member of VLAN 1 (the default VLAN).
   ii. You configure port 5 as an 802.1X authenticator port.
   iii. You configure port 5 to use an Authorized-Client VLAN.

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

■ If you expect friendly clients to connect without having 802.1X supplicant software running, provide a server on the Unauthorized-Client VLAN for downloading 802.1X supplicant software to the client, and a procedure by which the client initiates the download.

■ A client must either have a valid IP address configured before connecting to the switch, or download one through the Unauthorized-Client VLAN from a DHCP server. In the latter case, you will need to provide DHCP services on the Unauthorized-Client VLAN.

■ Ensure that the switch is connected to a RADIUS server configured to support authentication requests from clients using ports configured as 802.1X authenticators. (The RADIUS server should not be on the Unauthorized-Client VLAN.)

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

Caution

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

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

1. Enable 802.1X authentication on the individual ports you want to serve as authenticators. (The switch automatically disables LACP on the ports on which you enable 802.1X.) On the ports you will use as authenticators with VLAN operation, ensure that the (default) port-control parameter is set to **auto**. (Refer to "1. Enable 802.1X Authentication on Selected Ports" on page 10-14.) This setting requires a client to support 802.1X authentication (with 802.1X supplicant operation) and to provide valid credentials to get network access.

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

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

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

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

Determines the type of RADIUS authentication to use.

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

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

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

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

Syntax: radius host < *ip-address* >

Adds a server to the RADIUS configuration.

[key < *server-specific key-string* >]

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

radius-server key < *global key-string* >

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

4. Activate authentication on the switch.

Syntax: aaa port-access authenticator active

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

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

Note   If you want to implement the optional port security feature on the switch, you should first ensure that the ports you have configured as 802.1X authenticators operate as expected. Then refer to "Option For Authenticator Ports: Configure Port-Security To Allow Only 802.1X Devices" on page 10-36.

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

Configuring 802.1X Open VLAN Mode. Use these commands to actually configure Open VLAN mode. For a listing of the steps needed to prepare the switch for using Open VLAN mode, refer to "Preparation" on page 10-30.

Syntax: aaa port-access authenticator [e] < *port-list* >

[auth-vid < *vlan-id* >]

Configures an existing, static VLAN to be the Authorized-Client VLAN.

[< unauth-vid < *vlan-id* >]

Configures an existing, static VLAN to be the Unauthorized-Client VLAN.

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

- These two static VLANs already exist on the switch:
  - Unauthorized, VID = 80
  - Authorized, VID = 81

- Your RADIUS server has an IP address of 10.28.127.101. The server uses **rad4all** as a server-specific key string. The server is connected to a port on the Default VLAN.

- The switch's default VLAN is already configured with an IP address of 10.28.127.100 and a network mask of 255.255.255.0

```
HP Switch(config)# aaa authentication port-access eap-radius
```
Configures the switch for 802.1X authentication using an EAP-RADIUS server.

```
HP Switch(config)# aaa port-access authenticator 10-20
```
Configures ports 10 - 20 as 802.1 authenticator ports.

```
HP Switch(config)# radius host 10.28.127.101 key rad4all
```
Configures the switch to look for a RADIUS server with an IP address of 10.28.127.101 and an encryption key of rad4all.

```
HP Switch(config)# aaa port-access authenticator e 10-20 unauth-vid 80
```
Configures ports 10 - 20 to use VLAN 80 as the Unauthorized-Client VLAN.

```
HP Switch(config)# aaa port-access authenticator e 10-20 auth-vid 81
```
Configures ports 10 - 20 to use VLAN 81 as the Authorized-Client VLAN.

```
HP Switch(config)# aaa port-access authenticator active
```
Activates 802.1X port-access on ports you have configured as authenticators.

Inspecting 802.1X Open VLAN Mode Operation. For information and an example on viewing current Open VLAN mode operation, refer to "Viewing 802.1X Open VLAN Mode Status" on page 10-48.

802.1X Open VLAN Operating Notes

- Although you can configure Open VLAN mode to use the same VLAN for both the Unauthorized-Client VLAN and the Authorized-Client VLAN, this is *not* recommended. Using the same VLAN for both purposes allows unauthenticated clients access to a VLAN intended only for authenticated clients, which poses a security breach.

■ While an Unauthorized-Client VLAN is in use on a port, the switch temporarily removes the port from any other statically configured VLAN for which that port is configured as a member. Note that the Menu interface will still display the port's statically configured VLAN(s).

■ A VLAN used as the Unauthorized-Client VLAN should not allow access to resources that must be protected from unauthenticated clients.

■ If a port is configured as a tagged member of VLAN "X" that is not used as an Unauthorized-Client, Authorized-Client, or RADIUS-assigned VLAN, then the port returns to tagged membership in VLAN "X" upon successful client authentication. This happens even if the RADIUS server assigns the port to another, authorized VLAN "Y". Note that if RADIUS assigns VLAN "X" as an authorized VLAN, then the port becomes an untagged member of VLAN "X" for the duration of the client connection. After the client disconnects, the port returns to tagged membership in VLAN "X". (If there is no Authorized-Client or RADIUS-assigned VLAN, then an authenticated client without tagged VLAN capability can access only a statically configured, untagged VLAN on that port.)

■ When a client's authentication attempt on an Unauthorized-Client VLAN fails, the port remains a member of the Unauthorized-Client VLAN until the client disconnects from the port.

■ During an authentication session on a port in 802.1X Open VLAN mode, if RADIUS specifies membership in an untagged VLAN, this assignment overrides port membership in the Authorized-Client VLAN. If there is no Authorized-Client VLAN configured, then the RADIUS assignment overrides any untagged VLAN for which the port is statically configured.

■ If the only authenticated client on a port loses authentication during a session in 802.1X Open VLAN mode, the port VLAN membership reverts back to the Unauthorized-Client VLAN. If there is no Unauthorized-Client VLAN configured, then the client loses access to the port until it can reauthenticate itself. If there are multiple clients authenticated on the port, if one client loses access and attempts to re-authenticate, that client will be handled as a new client on the port.

■ The first client to authenticate on a port configured to support multiple clients will determine the port's VLAN membership for any subsequent clients that authenticate while an active session is already in effect.

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

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

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

> learn-mode port-access
>
> > *Configures port-security on the specified port(s) to allow only the first 802.1X-aware device the port detects.*
>
> action < none | send-alarm | send-disable >
>
> > *Configures the port's response (in addition to blocking unauthorized traffic) to detecting an intruder.*

Note   Port-Security operates with 802.1X authentication as described above only if the selected ports are configured as 802.1X; that is with the **control** mode in the port-access authenticator command set to **auto**. For example, to configure port 5 for 802.1X authenticator operation and display the result:

```
HP Switch(config)# aaa port-access authenticator 5 control auto
HP Switch(config)# show port-access authenticator 5 config
```

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

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

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

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

aaa port-access authenticator < *port-list* >

Enables 802.1X authentication on the port.

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

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

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

aaa port-access authenticator < *port-list* > control auto

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

Port Security

N o t e   If 802.1X port-access is configured on a given port, then port-security **learn-mode** for that port must be set to either **continuous** (the default) or **port-access**.

In addition to the above, to use port-security on an authenticator port, use the per-port **client-limit** option to control how many MAC addresses of 802.1X-authenticated devices the port is allowed to learn. (Using **client-limit** sets 802.1X to client-based operation on the specified ports.) When this limit is reached, no further devices can be authenticated until a currently authenticated device disconnects and the current delay period or logoff period has expired.

Configure the port access type.

Syntax: aaa port-access auth < *port-list* > client-limit < 1 - 2>

Configures client-based 802.1X authentication on the specified ports and sets the number of authenticated devices the port is allowed to learn. For more on this command, refer to "Configuring Switch Ports as 802.1X Authenticators" on page 10-14.)

— Or —

no aaa port-access auth < *port-list* > client-limit

> *Configures port-based 802.1X authentication on the specified ports, which opens the port.*

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

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. (A port can operate as both an authenticator and a supplicant.)

Example

Suppose that you want to connect two switches, where:

■   Switch "A" has port 1 configured for 802.1X supplicant operation

■   You want to connect port 1 on switch "A" to port 5 on switch "B".
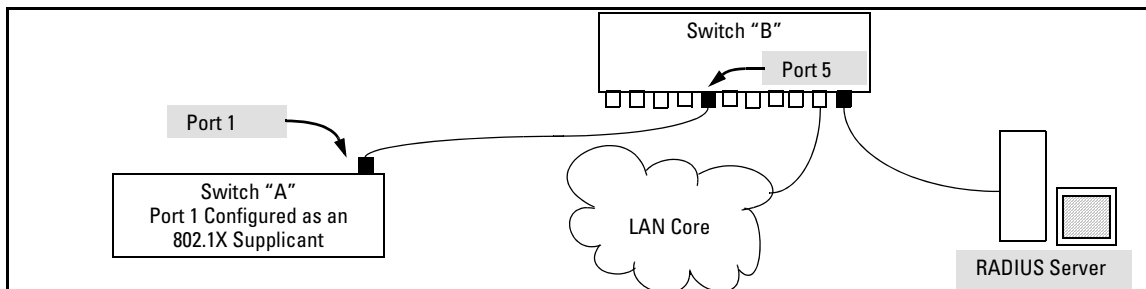
Figure 10-8. Example of Supplicant Operation

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

   - If, after the supplicant port sends the configured number of start request packets, it does not receive a response, it assumes that switch "B" is not 802.1X-aware, and transitions to the authenticated state. If switch "B" is operating properly and is not 802.1X-aware, then the link should begin functioning normally, but without 802.1X security.

   - If, after sending one or more start request packets, port 1 receives a request packet from port 5, then switch "B" is operating as an 802.1X authenticator. The supplicant port then sends a response/ID packet. If switch "B" is configured for RADIUS authentication, it forwards this request to a RADIUS server. If switch "B" is configured for Local 802.1X authentication (page 10-19), the authenticator compares the switch "A" response to its local username and password.

2. The RADIUS server then responds with an access challenge that switch "B" forwards to port 1 on switch "A".

3. Port 1 replies with a hash response based on its unique credentials. Switch "B" forwards this response to the RADIUS server.

4. The RADIUS server then analyzes the response and sends either a "success" or "failure" packet back through switch "B" to port 1.

   - A "success" response unblocks port 5 to normal traffic from port 1.

   - A "failure" response continues the block on port 5 and causes port 1 to wait for the "held-time" period before trying again to achieve authentication through port 5.

Supplicant Port Configuration

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 < *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 suppli-
cant 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 < *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 oper-*
> *tion. (Use one instance of the command for each*
> *parameter you want to configure The* **no** *form disables*
> *supplicant operation on the designated port(s).*

[identity < *username* >]

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

[secret]

> **Enter secret:** < *password* >
> **Repeat secret:** < *password* >
>
> *Sets the secret password to be used by the port suppli-*
> *cant when an MD5 authentication request is received*
> *from an authenticator. The switch prompts you to enter*
> *the secret password after the command is invoked.*

[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 10-39 for a description of how the port reacts to the authenticator response. (Default: 3).

[held-period < 0 - 65535 >]

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

[start-period < 1 - 300 >]

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

[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

Show Commands for Port-Access Authenticator

Syntax: show port-access authenticator [*port-list*]
[config | statistics | session-counters | vlan | clients]]

> *If you enter the* **show port-access authenticator** *command with-out an optional value, the following configuration informa-tion is displayed for all switch ports, or specified ports, that are enabled for 802.1X port-access authentication:*
>
> • **Port -access authenticator activated***: Are any switch ports configured to operate as 802.1X authenticators using the* **aaa port-access authenticator** *command?* **Yes** *or* **No**
>
> • **Allow RADIUS-assigned dynamic (GVRP) VLANs***: Are RADIUS-assigned dynamic (GVRP-learned) VLANs supported for authenticated and unauthenticated client sessions on the switch?* **Yes** *or* **No**
>
> • **Auth Clients***: Number of authorized clients*
>
> • **Unauth Clients***: Number of unauthorized clients*

Syntax: show port-access authenticator [*port-list*]
[config | statistics | session-counters | vlan | clients]| detailed]
—***Continued***—

> • **Untagged VLAN***: VLAN ID number of the untagged VLAN used in client sessions. If the switch supports MAC-based (untagged) VLANs,* **MACbased** *is displayed to show that multiple untagged VLANs are configured for authentication sessions.*
>
> • **Tagged VLANs***: Are tagged VLANs (statically configured or RADIUS-assigned) used for authenticated clients?* **Yes** *or* **No**

- **Port COS***:*
  Yes - *Client-specific CoS (Class of Service) values are applied to more than one authenticated client on the port.*
  No *- No client-specific CoS values are applied to any authenticated client on the port.*
  *<**cos-value**>- Numerical value of the CoS (802.1p priority) applied to inbound traffic from one authenticated client. For client-specific per-port CoS values, enter the* **show port-access web-based clients detailed** *command.*
- **Cntrl Dir***: Directions in which flow of incoming and outgoing traffic is blocked on 802.1X-aware port that has not yet entered the authenticated state:*
  Both: *Incoming and outgoing traffic is blocked on port until authentication occurs.*
  In: *Only incoming traffic is blocked on port before authentication occurs. Outgoing traffic with unknown destination addresses is flooded on the unauthenticated 802.1X-aware port.*

Information on ports not enabled for 802.1X port-access authentication is not displayed.

```
HP Switch(config)# show port-access authenticator 2-3
 Port Access Authenticator Status

  Port-access authenticator activated [No] : No
  Allow RADIUS-assigned dynamic (GVRP) VLANs [No] : No

        Auth      Unauth    Untagged Tagged             Cntrl
  Port Clients   Clients   VLAN     VLANs  Port COS   Dir
  ---- --------  --------  -------- ------ ---------  -------
  2    1         0         1               7          In
  3    1         0         1               5          In

```

Figure 10-9. Example of show port-access authenticator Command

The information displayed with the **show port-access authenticator** command for individual (config | statistics | session-counters | vlan | clients) options is described below.

```
HP Switch(config)# show port-access authenticator config

 Port Access Authenticator Configuration

Port-access authenticator activated [No] : No

       | Re-auth Access    Max   Quiet   TX        Supplicant Server
  Port | Period  Control   Reqs  Period  Timeout   Timeout    Timeout
  ---- + ------- --------  ----- ------- --------  ---------- --------
  1    | No      Auto      2     60      30        30         30
  2    | No      Auto      2     60      30        30         30
```

Figure 10-10.Example of show port-access authenticator config Command

Table 10-2. Field Descriptions of show port-access authenticator config Command Output (Figure 10-10)

Field	Description
Port-access authenticator activated	Whether 802.1X authentication is enabled or disabled on specified port(s).
Port	Port number on switch.
Re-auth Period	Period of time (in seconds) after which clients connected to the port need to be re-authenticated.
Access Control	Port's authentication mode: **Auto:** Network access is allowed to any connected device that supports 802.1X authentication and provides valid 802.1X credentials. **Authorized:** Network access is allowed to any device connected to the port, regardless of whether it meets 802.1X criteria. **Unauthorized:** Network access is blocked to any device connected to the port, regardless of whether the device meets 802.1X criteria.
Max reqs	Number of authentication attempts that must time-out before authentication fails and the authentication session ends.
Quiet Period	Period of time (in seconds) during which the port does not try to acquire a supplicant.
TX Timeout	Period of time (in seconds) that the port waits to retransmit the next EAPOL PDU during an authentication session.
Supplicant Timeout	Period of time (in seconds) that the switch waits for a supplicant response to an EAP request.
Server Timeout	Period of time (in seconds) that the switch waits for a server response to an authentication request.

Syntax: show port-access authenticator statistics [*port-list*]

> *Displays statistical information for all switch ports or spec-ified ports that are enabled as 802.1X authenticators, includ-ing:*
>
> - *Whether port-access authentication is enabled*
> - *Whether RADIUS-assigned dynamic VLANs are supported*
> - *802.1X supplicant's MAC address as determined by the content of the last EAPOL frame received on the port*
> - *802.1X traffic statistics from received and transmitted packets*
>
> *802.1X configuration information for ports that are not enabled as an 802.1X authenticators is not displayed.*

```
HP Switch(config)# show port-access authenticator statistics

 Port Access Authenticator Statistics

  Port-access authenticator activated [No] : Yes
  Allow RADIUS-assigned dynamic (GVRP) VLANs [No] : No

       Source         TX     TX    RX     RX      RX      RX    RX
  Port MAC address    ReqId  Req   Start  Logoff  RespId  Resp  Errors
  ---- -------------  ------ ----- ------ ------- ------- ----- -------
  2    001560-b3ea48  1      0     0      0       0       0     0
```

Figure 10-11. Example of show port-access authenticator statistics Command

Syntax: show port-access authenticator session-counters [*port-list*]

> *Displays information for active 802.1X authentication ses-sions on all switch ports or specified ports that are enabled as 802.1X authenticators, including:*
>
> - *802.1X frames received and transmitted on each port*
> - *Duration and status of active 802.1X authentication sessions (**in-progress** or **terminated**)*
> - *User name of 802.1X supplicant included in 802.1X response packets, configured with the **aaa port-access supplicant identity** <username> command (see page 13-49).*
>
> *802.1X configuration information for ports that are not enabled as an 802.1X authenticators is not displayed.*

```
HP Switch(config)# show port-access authenticator session-counters

 Port Access Authenticator Session Counters

  Port-access authenticator activated [No] : Yes
  Allow RADIUS-assigned dynamic (GVRP) VLANs [No] : No

                                  Session     Session
  Port Frames In  Frames Out  Time(sec.) Status      User
  ---- ---------- ----------- ----------- ----------- -------
   2    45623      45623         20         in-progress bert
```

Figure 10-12. Example of show port-access authenticator session-counters Command

Syntax: show port-access authenticator vlan [*port-list*]

> *Displays the following information on the VLANs configured for use in 802.1X port-access authentication on all switch ports, or specified ports, that are enabled as 802.1X authenticator:*
>
> • *Authentication mode used on each port, configured with the* **aaa port-access authenticator control** *command (see page 10-17)*
> • *VLAN ID (if any) to be used for traffic from 802.1X-authenticated clients*
> • *VLAN ID (if any) to be used for traffic from unauthenticated clients*
>
> *802.1X configuration information for ports that are not enabled as an 802.1X authenticators is not displayed.*

```
HP Switch(config)# show port-access authenticator vlan

 Port Access Authenticator VLAN Configuration

  Port-access authenticator activated [No] : Yes
  Allow RADIUS-assigned dynamic (GVRP) VLANs [No] : No

        Access    Unauth   Auth
  Port Control   VLAN ID  VLAN ID
  ---- --------  -------- --------
   2    Auto       0        0
```

Figure 10-13.Example of show port-access authenticator vlan Command

Syntax: show port-access authenticator clients [*port-list*]

> *Displays the name, and address for each 802.1X port-access-authenticated client on the switch. One or two authenticated clients may be displayed for the same port.*
> - *If an 802.1X-authenticated client uses an IPv6 address,* **n/a - IPv6** *is displayed.*

```
HP Switch(config)# show port-access authenticator clients

 Port Access Authenticator Client Status

Port  Client Name  MAC Address    IP Address
-----　------------　-------------　--------------
1     webuser1     001321-eb8063  192.192.192.192
1     webuser2     001560-b3ea48  n/a - no info
3     webuser4     000000-111112  n/a
```

Figure 10-14. Example of show port-access authenticator clients Command Output

Syntax: show port-access authenticator clients <*port-list*> detailed

> *Displays detailed information on the status of 802.1X-authenticated client sessions on specified ports.*

```
HP Switch(config)# show port-access authenticator clients 5 detailed

 Port Access Authenticator Client Status Detailed

  Client Base Details :
   Port         : 5
   Session Status : Open          Session Time(sec) : 999999999
   Frames In    : 999999999       Frames Out        : 99999999
   Username     : webuser1        MAC Address       : 001321-eb8063
   IP           : 2001:fecd:ba23:cd1f:dcb1:1010:9234:4088

  Access Policy Details :
   COS Map      : 70000000
   Untagged VLAN : 3096
   Tagged VLANs : 1, 3, 5, 6, 334, 2066
```

Figure 10-15. Example of show port-access authenticator clients detailed Command Output

Viewing 802.1X Open VLAN Mode Status

You can examine the switch's current VLAN status by using the **show port-access authenticator** and **show vlan** < *vlan-id* > commands as illustrated in this section. Figure 10-16 shows an example of **show port-access authenticator** output, and table 10-3 describes the data that this command displays. Figure 10-17 shows related VLAN data that can help you to see how the switch is using statically configured VLANs to support 802.1X operation.

```
HP Switch(config)# show port-access authenticator 1-4        An Unauth VLAN ID appearing in the
                                                              Current VLAN ID column for the same
 Port Access Authenticator Status                            port indicates an unauthenticated client
                                                              is connected to this port.
  Port-access authenticator activated [No] : Yes             (Assumes that the port is not a statically
                                                              configured member of VLAN 100.)

            Access   Authenticator   Authenticator  Unauth    Auth     Current
 Port Status Control State           Backend State  VLAN ID   VLAN ID  VLAN ID
 ---- ------ -------- -------------- -------------- --------  -------- -------
 1    Closed Auto     Connecting     Idle            100       101      100

 2    Open   Auto     Authorized     Idle            100       101      101

 3    Closed Auto     Connecting     Idle            100       0        100

 4    Closed Auto     Disconnected   Idle            100       101      No PVID
```

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

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

4. A "0" in the row for port 3 indicates there is no Authorized VLAN configured for port 3.
5. "No PVID" means there is currently no untagged VLAN membership on port 4.

Figure 10-16. Example Showing Ports Configured for Open VLAN Mode

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

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

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

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

Table 10-3.   Open VLAN Mode Status

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

Table 10-3. Open VLAN Mode Status

Status Indicator	Meaning
Unauthorized VLAN ID	< *vlan-id* >: Lists the VID of the static VLAN configured as the unauthorized VLAN for the indicated port.
	0: No unauthorized VLAN has been configured for the indicated port.
Authorized VLAN ID	< *vlan-id* >: Lists the VID of the static VLAN configured as the authorized VLAN for the indicated port.
	0: No authorized VLAN has been configured for the indicated port.
Current VLAN ID	< *vlan-id* >: Lists the VID of the static, untagged VLAN to which the port currently belongs.
	No PVID: The port is not an untagged member of any VLAN.

Syntax: show vlan < *vlan-id* >

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

```
HP Switch(config)# show vlan 1

 Status and Counters - VLAN Information - VLAN 1

 VLAN ID : 1
 Name : DEFAULT_VLAN
 Status : Static  Voice : No

 Port Information Mode     Unknown VLAN Status
 ---------------- -------- ------------ ----------
 1                Untagged Learn        Up
 3                Untagged Learn        Up
 4                Untagged Learn        Down
 6                Untagged Learn        Down
 7                Untagged Learn        Up
 8                Untagged Learn        Down
 9                Untagged Learn        Down
 10               Untagged Learn        Down

 Overridden Port VLAN configuration


 Port Mode
 ---- ------------
 2    Untagged
 5    Untagged
```

Note that ports 2 and 5 are not in the upper listing, but are included under "Overridden Port VLAN configuration". This shows that static, untagged VLAN memberships on ports 2 and 5 have been overridden by temporary assignment to the authorized or unauthorized VLAN. Using the **show port-access authenticator < *port-list* >** command shown in figure 10-16 provides details.

Figure 10-17. Example of Showing a VLAN with Ports Configured for Open VLAN Mode

Show Commands for Port-Access Supplicant

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

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

> *Shows the port-access supplicant configuration (excluding the **secret** parameter) for all ports or < port-list> ports configured on the switch as supplicants. The Supplicant State can include the following:*
> **Connecting -** *Starting authentication.*
> **Authenticated -** *Authentication completed (regardless of whether the attempt was successful).*
> **Acquired -** *The port received a request for identification from an authenticator.*
> **Authenticating -** *Authentication is in progress.*
> **Held -** *Authenticator sent notice of failure. The supplicant port is waiting for the authenticator's held-period (page 10-40).*
>
> *For descriptions of the supplicant parameters, refer to "Configuring a Supplicant Switch Port" on page 10-40.*

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

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

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

■   The supplicant port detects a different authenticator device.

■   You use the **aaa port-access supplicant < *port-list* > clear-statistics** command to clear the statistics for the supplicant port.

■   The switch reboots.

Thus, if the supplicant's link to the authenticator fails, the supplicant retains the transaction statistics it most recently received until one of the above events occurs. Also, if you move a link with an authenticator from one

supplicant port to another without clearing the statistics data from the first port, the authenticator's MAC address will appear in the supplicant statistics for both ports.

How RADIUS/802.1X Authentication Affects VLAN Operation

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

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

VLAN Assignment on a Port

Following client authentication, VLAN configurations on a port are managed as follows when you use 802.1X, MAC, or Web authentication:

■ The port resumes membership in any tagged VLANs for which it is already assigned in the switch configuration. Tagged VLAN membership allows a port to be a member of multiple VLANs simultaneously.

■ The port is temporarily assigned as a member of an untagged (static or dynamic) VLAN for use during the client session according to the following order of options.

  a. The port joins the VLAN to which it has been assigned by a RADIUS server during client authentication.

b.   If RADIUS authentication does not include assigning the port to a VLAN, then the switch assigns the port to the authorized-client VLAN configured for the authentication method.

c.   If the port does not have an authorized-client VLAN configured, but is configured for membership in an untagged VLAN, the switch assigns the port to this untagged VLAN.

Operating Notes

■   During client authentication, a port assigned to a VLAN by a RADIUS server or an authorized-client VLAN configuration is an untagged member of the VLAN for the duration of the authenticated session. This applies even if the port is also configured in the switch as a tagged member of the same VLAN. The following restrictions apply:

•   If the port is assigned as a member of an untagged *static* VLAN, the VLAN must already be configured on the switch. If the static VLAN configuration does not exist, the authentication fails.

•   If the port is assigned as a member of an untagged *dynamic* VLAN that was learned through GVRP, the dynamic VLAN configuration must exist on the switch at the time of authentication and GVRP-learned dynamic VLANs for port-access authentication must be enabled.

    If the dynamic VLAN does not exist or if you have not enabled the use of a dynamic VLAN for authentication sessions on the switch, the authentication fails.

■   To enable the use of a GVRP-learned (dynamic) VLAN as the untagged VLAN used in an authentication session, enter the **aaa port-access gvrp-vlans** command, as described in "Enabling the Use of GVRP-Learned Dynamic VLANs in Authentication Sessions" on page 10-57.

■   Enabling the use of dynamic VLANs in an authentication session offers the following benefits:

•   You avoid the need of having static VLANs pre-configured on the switch.

•   You can centralize the administration of user accounts (including user VLAN IDs) on a RADIUS server.

For information on how to enable the switch to dynamically create 802.1Q-compliant VLANs on links to other devices using the GARP VLAN Registration Protocol (GVRP), see the chapter on "GVRP" in the *Advanced Traffic Management Guide*.

■ For an authentication session to proceed, a port must be an untagged member of the (static or dynamic) VLAN assigned by the RADIUS server (or an authorized-client VLAN configuration). The port temporarily drops any current untagged VLAN membership.

If the port is not already a member of the RADIUS-assigned (static or dynamic) untagged VLAN, the switch temporarily reassigns the port as an untagged member of the required VLAN (for the duration of the session). *At the same time, if the port is already configured as an untagged member of a different VLAN, the port loses access to the other VLAN for the duration of the session.* (A port can be an untagged member of only one VLAN at a time.)

When the authentication session ends, the switch removes the temporary untagged VLAN assignment and re-activates the temporarily disabled, untagged VLAN assignment.

■ If GVRP is already enabled on the switch, the temporary untagged (static or dynamic) VLAN created on the port for the authentication session is advertised as an existing VLAN.

If this temporary VLAN assignment causes the switch to disable a different untagged static or dynamic VLAN configured on the port (as described in the preceding bullet and in "Example of Untagged VLAN Assignment in a RADIUS-Based Authentication Session" on page 10-55), the disabled VLAN assignment is not advertised. When the authentication session ends, the switch:

• Removes the temporary untagged VLAN assignment and stops advertising it.

• Re-activates and resumes advertising the temporarily disabled, untagged VLAN assignment.

■ If you modify a VLAN ID configuration on a port during an 802.1X, MAC, or Web authentication session, the changes do not take effect until the session ends.

■ When a switch port is configured with RADIUS-based authentication to accept multiple 802.1X and/or MAC or Web authentication client sessions, all authenticated clients must use the same port-based, untagged VLAN membership assigned for the earliest, currently active client session.

Therefore, on a port where one or more authenticated client sessions are already running, all such clients are on the same untagged VLAN. If a RADIUS server subsequently authenticates a new client, but attempts to re-assign the port to a different, untagged VLAN than the one already in use for the previously existing, authenticated client sessions, the connection for the new client will fail.

Example of Untagged VLAN Assignment in a RADIUS-Based Authentication Session

The following example shows how an untagged static VLAN is temporarily assigned to a port for use during an 802.1X authentication session. In the example, an 802.1X-aware client on port 2 has been authenticated by a RADIUS server for access to VLAN 22. However, port 2 is not configured as a member of VLAN 22 but as a member of untagged VLAN 33 as shown in Figure 10-18.

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

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

  Port    default_vlan     vlan_22          vlan_33         vlan_44
  ---- +  -----------     -----------     -----------     -----------
   1   |  Untagged         Tagged          No_ _ _ _       No
   2   |  No               No              Untagged        No
   3   |  Untagged         Forbid          Forbid          Forbid
   4   |  Untagged         Tagged          Tagged          Tagged
   .   |     .                .               .               .
   .   |     .                .               .               .

  Actions->    Cancel      Edit      Save      Help

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

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

Figure 10-18. Example of an Active VLAN Configuration

In figure 10-18, 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).

To view the temporary VLAN assignment as a change in the active configuration, use the **show vlan <*vlan-id*>** command as shown in Figure 10-17 where <*vlan-id*> is the (static or dynamic) VLAN used in the authenticated client session.

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

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

  Port  Mode
  ----  ----
  2     No
```

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

Note: With the current VLAN configuration (figure 10-18), the only time port 2 appears in this **show vlan 22** listing is during an 802.1X session with an attached client. Otherwise, port 2 is not listed.

Figure 10-19. The Active Configuration for VLAN 22 Temporarily Changes for the 802.1X Session

However, as shown in Figure 10-18, because VLAN 33 is configured as untagged on port 2 and because a port can be untagged on only one VLAN, port 2 loses access to VLAN 33 for the duration of the 802.1X session on VLAN 22.

You can verify the temporary loss of access to VLAN 33 by entering the **show vlan 33** command as shown in Figure 10-20.

Even though port 2 is configured as Untagged on (static) VLAN 33 (see figure 10-18), 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 2 to VLAN 33.

```
HP Switch(config)# 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
  --------------- -------- ------------ ----------
  4               Tagged   Learn        Up

  Overridden Port VLAN configuration
  Port  Mode
  ----  ----
  2     Untagged
```

Figure 10-20. The Active Configuration for VLAN 33 Temporarily Drops Port 2 for the 802.1X Session

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

```
             HP Switch(config)# 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
              --------------- --------  ------------ ----------
              2                         Untagged Learn       Down
              4                         Tagged   Learn       Down
```

After the 802.1X session on VLAN 22 ends, the active configuration again includes VLAN 33 on port 2.
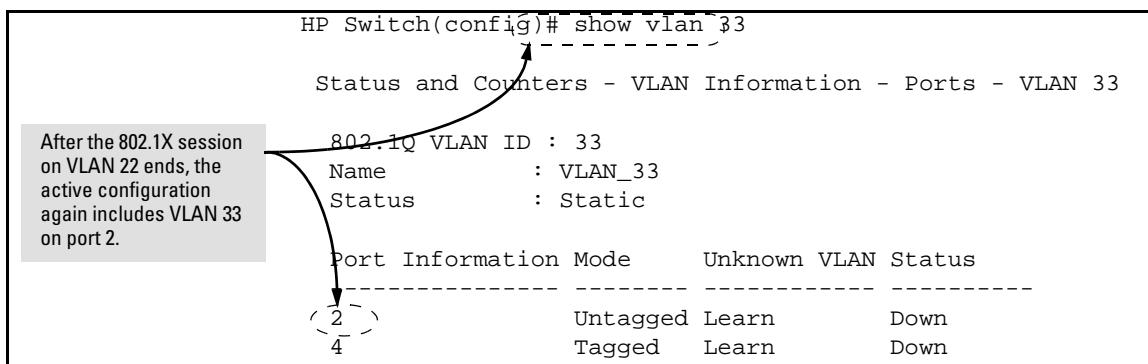
Figure 10-21. The Active Configuration for VLAN 33 Restores Port 2 After the 802.1X Session Ends

Enabling the Use of GVRP-Learned Dynamic VLANs in Authentication Sessions

Syntax: aaa port-access gvrp-vlans

Enables the use of dynamic VLANs (learned through GVRP) in the temporary untagged VLAN assigned by a RADIUS server on an authenticated port in an 802.1X, MAC, or Web authentication session.

*Enter the **no** form of this command to disable the use of GVRP-learned VLANs in an authentication session.*

For information on how to enable a switch to dynamically create 802.1Q-compliant VLANs, see the chapter on "GVRP" in the Advanced Traffic Management Guide.

Notes:
1. If a port is assigned as a member of an untagged dynamic VLAN, the dynamic VLAN configuration must exist at the time of authentication and GVRP for port-access authentication must be enabled on the switch.

If the dynamic VLAN does not exist or if you have not enabled the use of a dynamic VLAN for authentication sessions on the switch, the authentication fails.

Syntax: aaa port-access gvrp-vlans
—Continued—

2. After you enable dynamic VLAN assignment in an authentication session, it is recommended that you use the **interface unknown-vlans** *command on a per-port basis to prevent denial-of-service attacks. The* **interface unknown-vlans** c*ommand allows you to:*

- *Disable the port from sending advertisements of existing GVRP-created VLANs on the switch.*
- *Drop all GVRP advertisements received on the port.*

For more information, refer to the chapter on "GVRP" in the Advanced Traffic Management Guide.

3. If you disable the use of dynamic VLANs in an authentication session using the **no aaa port-access gvrp-vlans** *command, client sessions that were authenticated with a dynamic VLAN continue and are not deauthenticated. (This behavior differs form how static VLAN assignment is handled in an authentication session. If you remove the configuration of the static VLAN used to create a temporary client session, the 802.1X, MAC, or Web authenticated client is deauthenticated.)*

However, if a RADIUS-configured dynamic VLAN used for an authentication session is deleted from the switch through normal GVRP operation (for example, if no GVRP advertisements for the VLAN are received on any switch port), authenticated clients using this VLAN are deauthenticated.

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.

Configuring and Monitoring Port Security

Overview

Feature	Default	Menu	CLI	Web
Displaying Current Port Security	n/a	—	page 11-6	page 11-28
Configuring Port Security	disabled	—	page 11-9	page 11-28
Retention of Static Addresses	n/a	—	page 11-14	n/a
MAC Lockdown	disabled	—	page 11-19	
MAC Lockout	disabled	—	page 11-26	
Intrusion Alerts and Alert Flags	n/a	page 11-35	page 11-33	page 11-36

Port Security (Page 11-2). This feature enables you to 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.

N o t e

This feature does not prevent intruders from receiving broadcast and multicast traffic. Also, Port Security and MAC Lockdown are mutually exclusive on a switch. If one is enabled, then the other cannot be used.

MAC Lockdown (Page 11-19). This feature, also known as "Static Addressing", is used to prevent station movement and MAC address "hijacking" by allowing a given MAC address to use only an assigned port on the switch. MAC Lockdown also restricts the client device to a specific VLAN. (See also the **Note**, above.)

MAC Lockout (Page 11-26). This feature enables you to block a specific MAC address so that the switch drops all traffic to or from the specified address.

Port Security

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.

Eavesdrop Protection. Using either the port-security command or the switch's web browser interface to enable port security on a given port automatically enables eavesdrop prevention on 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 port security is configured, 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 IMC.

■ Alert Log entries in the switch's web browser interface

■ Event Log entries in the console interface

■ Intrusion Log entries in the menu interface, CLI, or web browser interface

For any port, you can configure the following:

■ **Action:** Used when a port detects an intruder. Specifies whether to send an SNMP trap to a network management station and whether to disable the port.

■ **Address Limit:** Sets the number of authorized MAC addresses allowed on the port.

■ **Learn-Mode:** Specify how the port acquires authorized addresses.

  • **Continuous:** Allows the port to learn addresses from inbound traffic from any connected device. This is the default setting.

  • **Limited-Continuous:** Sets a finite limit (1 - 32) to the number of learned addresses allowed per port.

- **Static:** Enables you to set a fixed limit on the number of MAC addresses authorized for the port and to specify some or all of the authorized addresses. (If you specify only some of the authorized addresses, the port learns the remaining authorized addresses from the traffic it receives from connected devices.)

- **Configured:** Requires that you specify all MAC addresses authorized for the port. The port is not allowed to learn addresses from inbound traffic.

■ **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.)

■ **Port Access:** Allows only the MAC address of a device authenticated through the switch's 802.1X Port-Based access control. Refer to chapter 10, Configuring Port-Based and Client-Based Access Control (802.1X).

For configuration details, refer to "Configuring Port Security" on page 11-9.

Eavesdrop Protection

Configuring port security on a given switch port automatically enables eavesdrop protection for that port. This prevents use of the port to flood unicast packets addressed to MAC addresses unknown to the switch. This blocks unauthorized users from eavesdropping on traffic intended for addresses that have aged-out of the switch's address table. (Eavesdrop prevention does not affect multicast and broadcast traffic, meaning that the switch floods these two traffic types out a given port regardless of whether port security is enabled on that port.)

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:
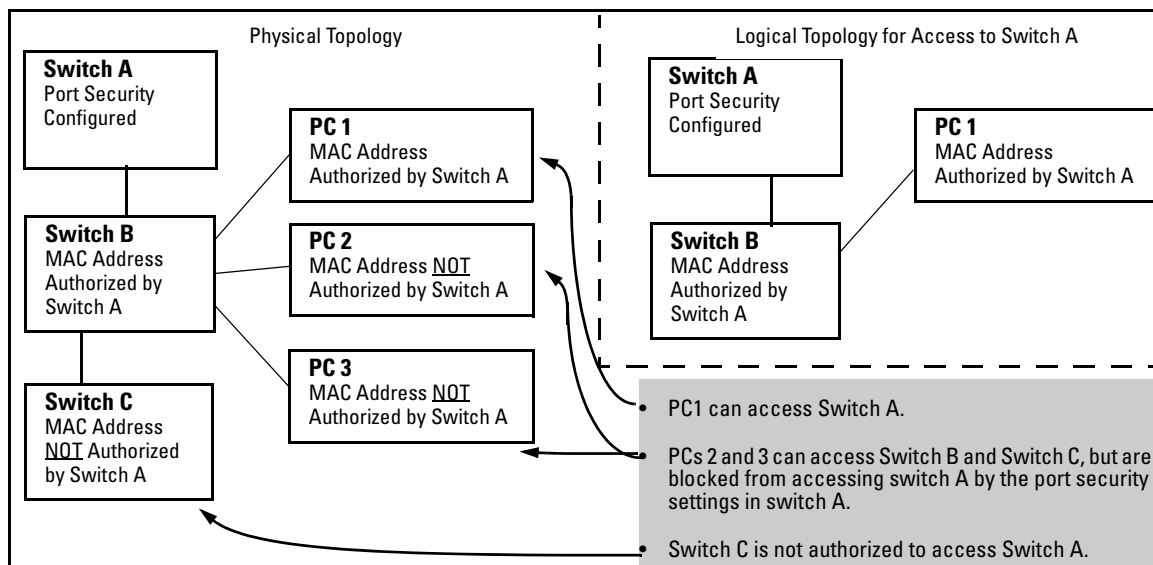
Figure 11-1. Example of How Port Security Controls Access

N o t e
Broadcast and Multicast traffic is always allowed, 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 port security?

    b. Which devices (MAC addresses) are authorized on each 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.

Port Security Command Options and Operation

Port Security Commands Used in This Section

show port-security	11-6
show mac-address	11-7
port-security	11-9
< port-list >	11-9
learn-mode	11-9
address-limit	11-12
mac-address	11-13
action	11-13
clear-intrusion-flag	11-14
no port-security	11-14

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

N o t e   Use the global configuration level to execute port-security configuration
commands.

Port Security Display Options

You can use the CLI to display the current port-security settings and to list the
currently authorized MAC addresses the switch detects on one or more ports.

Displaying Port Security Settings.

Syntax: show port-security
            show port-security *<port number>*
            show port-security [*<port number>*-*<port number>*]. . .[,*<port number>*]

> *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*
>
> *Without port parameters*, **show port-security** *displays Operating*
> *Control settings for all ports on a switch.*

```
HP Switch(config)# show port-security
 Port Security
  Port Learn Mode |         Action
  ---- ---------- + -----------------------
   1    Static     | Send Alarm, Disable Port
   2    Static     | Send Alarm, Disable Port
   3    Static     | Send Alarm
   4    Static     | Send Alarm
   5    Static     | Send Alarm
   6    Static     | Send Alarm
   7    Continuous | None
   8    Continuous | None
```

Figure 11-2. Example Port Security Listing (Ports 7 and 8 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 spec-
ified ports on a switch. The following example lists the full port security
configuration for a single port:

```
HP Switch(config)# show port-security 3
 Port Security
  Port :  3

  Learn Mode :  Static        Address Limit : 1

  Action : Send Alarm

  Authorized Addresses
  --------------------
  00906d-fdcc00
```

Figure 11-3. Example of the Port Security Configuration Display for a Single Port

The next 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:

```
HP Switch(config)# show port-security 1-3,6,8
```

Listing Authorized and Detected MAC Addresses.

Syntax: show mac-address [*port-list* | *mac-address* | vlan < *vid* >]

> *Without an optional parameter,* **show mac-address** *lists the authorized MAC addresses that the switch detects on all ports.*

> **mac-address:** *Lists the specified MAC address with the port on which it is detected as an authorized address.*
> **port list:** *Lists the authorized MAC addresses detected on the specified port(s).*
> **vlan < vid >:** *Lists the authorized MAC addresses detected on ports belonging to the specified VLAN.*

```
HP Switch(config)# show mac-address

 Status and Counters - Port Address Table

  MAC Address    Located on Port
  ------------- ---------------
  0004ea-B4d980 7
  0004ea-B4d9ee 7
  000a57-4d8d40 5
        .
        .
        .
  00a0c9-f1786f 5


HP Switch(config)# show mac-address 7
 Status and Counters - Port Address Table - 7
  MAC Address
  -------------
  0004ea-84d980
  0004ea-84d9ee


HP Switch(config)# show mac-address 000a57-4d8d40
 Status and Counters - Address Table 000a57-4d8d40
  MAC Address : 000a57-4d8d40
  Located on Port : 5


HP Switch(config)# show mac-address vlan 1
 Status and Counters - Address Table VLAN 1
  MAC Address    Located on Port
  ------------- ---------------
  0004ea-84d980 7
  0004ea-84d9ee 7
  000a57-4d8d40 6
        .
        .
        .
  00a0o9-f17B6f 5
```

Figure 11-4. Examples of Show Mac-Address Outputs

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

[e] <port-list>< learn-mode | address-limit | mac-address | action | clear-intrusion-flag >

*< **port-list >**: Specifies a list of one or more ports to which the port-security command applies.*

learn-mode < continuous | static | port-access | configured | limited-continuous >

For the specified port:
- *Identifies the method for acquiring authorized addresses.*
- *On switches covered in this guide, automatically invokes eavesdrop protection. (Refer to "Eavesdrop Protection" on page 11-3.)*

continuous *(Default): Appears in the factory-default setting or when you execute* **no port-security.** *Allows the port to learn addresses from the 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 in the learn continuous mode will "age out" and be automatically deleted if they are not used regularly. The default age time is five minutes.*

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. You can set the MAC age out time using the CLI, SNMP, Web, or menu interfaces. For more information on the* **mac-age-time** *command refer to the chapter titled "Interface Access and System Information" in the* Basic Operation Guide *for your switch.*

— Continued —

Syntax: port-security *(Continued)*

learn-mode < continuous | static | port-access | configured | limited-
continuous > *(Continued)*

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 (explained below) 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:

You use **mac-address** *to authorize MAC address 0060b0-880a80 for port 4.*

You use **address-limit** *to allow three devices on port 4 and the port detects these MAC addresses:*

1. 080090-1362f2

2. 00f031-423fc1

3. 080071-0c45a1

4. 0060b0-880a80 *(the address you authorized with the* **mac-address** *parameter)*

In this example port 4 would assume the following list of authorized addresses:

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)*

The remaining MAC address detected by the port, 080071-0c45a1, *is not allowed and is handled as an intruder. Learned addresses that become authorized do* ***not*** *age-out. See also "Retention of Static Addresses" on page 11-14.*

— Continued —

Syntax: port-security *(Continued)*

learn-mode < continuous | static | port-access | configured | limited-continuous > *(Continued)*

> *Caution: Using the* **static** *parameter with a device limit greater than the number of MAC addresses specified with* **mac-address** *can allow an unwanted device to become "authorized". This is because the port, to fulfill the number of devices allowed by the* **address-limit** *parameter (se below), automatically adds devices it detects until it reaches the specified limit.*

> *Note: If 802.1X port-access is configured on a given port, then port-security learn-mode must be set to either* **continuous** *(the default) or* **port-access***.*

port-access*: Enables you to use Port Security with (802.1X) Port-Based Access Control. Refer to chapter 10, Configuring Port-Based and Client-Based Access Control (802.1X).*

configured*: Must specify which MAC addresses are allowed for this port. Range is 1 (default) to 8 and addresses are not ageable. Addresses are saved across reboots.*

limited-continuous*: Also known as MAC Secure, or "limited" mode. The limited parameter sets a finite limit to the number of learned addresses allowed per port. (You can set the range from 1, the default, to a maximum of 32 MAC addresses which may be learned by each port.)*

All addresses are ageable, meaning they are automatically removed from the authorized address list for that port after a certain amount of time. Limited mode and the address limit are saved across reboots, but addresses which had been learned are lost during the reboot process.

Addresses learned in the limited mode are normal addresses learned from the network until the limit is reached, but they are not configurable. (You cannot enter or remove these addresses manually if you are using **learn-mode** *with the* **limited-continuous** *option.)*

—Continued—

Syntax: port-security *(Continued)*

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. You can set the MAC age out time using the CLI, SNMP, Web, or menu interfaces. For more on the* **mac-age-time** *command, refer to the chapter titled "Interface Access and System Information" in the* Basic Operation Guide *for your switch. To set the learn-mode to* **limited** *use this command syntax:*

port-security <*port-list*> learn-mode limited address-limit < 1..32 > action < none | send-alarm | send-disable >

The default address-limit is **1** *but may be set for each port to learn up to 32 addresses. The default action is* **none**. *To see the list of learned addresses for a port use the command:*

show mac < *port-list* >

address-limit < *integer* >

When **learn-mode** *is set to* **static**, **configured**, *or* **limited-continuous**, *the* **address-limit** *parameter specifies how many authorized devices (MAC addresses) to allow. Range: 1 (the default) to 8 for static and configured modes. For* **learn-mode** *with the* **limited-continuous** *option, the range is 1-32 addresses.*

—Continued—

Syntax: port-security *(Continued)*

mac-address [<*mac-addr*>] [<*mac-addr*>] . . . [<*mac-addr*>]

Available for **learn-mode** *with the,* **static**, **configured**, *or* **limited-continuous** *option. Allows up to eight authorized devices (MAC addresses) per port, depending on the value specified in the* **address-limit** *parameter. The* **mac-address** **limited-continuous** *mode allows up to 32 authorized MAC addresses per port.*

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. Learned addresses that become authorized do **not** age-out. See also "Retention of Static Addresses" on page 11-14.*

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: *Prevents an SNMP trap from being sent.* **none** *is the default value.*

send-alarm*: Sends an intrusion alarm. Causes the switch to send an SNMP trap to a network management station.*

send-disable: *Sends alarm and disables the port. Available only in the* **static**, **port-access**, **configured**, *or* **limited learn** *modes. Causes the switch to send an SNMP trap to a network management station and disable the port. If you subsequently re-enable the port without clearing the port's intrusion flag, the port blocks further intruders, but the switch will not disable the port again until you reset the intrusion flag. See the Note on 11-31.*

For information on configuring the switch for SNMP management, refer to the Management and Configuration Guide *for your switch.*

—Continued—

Syntax: port-security *(Continued)*

> clear-intrusion-flag

>> *Clears the intrusion flag for a specific port. (See "Reading Intrusion Alerts and Resetting Alert Flags" on page 11-29.)*

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

>> *Removes the specified learned MAC address(es) from the specified port.*

Retention of Static Addresses

Static MAC addresses do not age-out. MAC addresses learned by using **learn-mode continuous** or **learn-mode limited-continuous** age out according to the currently configured MAC age time. (For information on the **mac-age-time** command, refer to the chapter titled "Interface Access and System Information" in the *Basic Operation Guide* for your switch.)

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:

■ 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.

To remove an address learned using either of the preceding methods, do one of the following:

■ 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.

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:

- 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.

Specifying Authorized Devices and Intrusion Responses. This example configures port 1 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.

```
HP Switch(config)# port-security 1 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.

```
HP Switch(config)# port-security 1 learn-mode static
mac-address 0c0090-123456 action send-disable
```

This example configures port 5 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, but allow the intruder access to the network.

```
HP Switch(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 1 allows two authorized devices, but has only one device in its Authorized Address list:

```
                        HP Switch(config)# show port-security 1
Although the Address     Port Security
Limit is set to 2, only
one device has been
authorized for this      Port : 1
port. In this case you   Learn Mode [Continuous] : Static       Address Limit [1] : 2
can add another          Action [None] : None
without having to also
increase the Address
Limit.                   Authorized Addresses
                         --------------------              The Address Limit has not
                         0c0090-123456                     been reached.
```

Figure 11-5. Example of Adding an Authorized Device to a Port

With the above configuration for port 1, the following command adds the 0c0090-456456 MAC address as the second authorized address.

```
HP Switch(config)# port-security 1 mac-address 0c0090-
456456
```

After executing the above command, the security configuration for port 1 would be:

```
HP Switch(config)# show port-security 1
 Port Security

  Port : 1
  Learn Mode [Continuous] : Static       Address Limit [1] : 2
  Action [None] : None

  Authorized Addresses
  --------------------              The Address Limit has been
  0c0090-123456                     reached.
  0c0090-456456
```

Figure 11-6. 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 subse-

quently 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 1 allows one authorized device and already has a device listed:

```
HP Switch(config)# show port-security 1
 Port Security
  Port : 1
  Learn Mode [Continuous] : Static      Address Limit [1] : 1
  Action [None] : None

  Authorized Addresses
  -------------------
  0c0090-123456
```

Figure 11-7. Example of Port Security on Port 1 with an Address Limit of "1"

To add a second authorized device to port 1, execute a **port-security** command for port 1 that raises the address limit to 2 and specifies the additional device's MAC address. For example:

```
HP Switch(config)# port-security 1 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". Refer to the command syntax listing under "Configuring Port Security" on page 11-9.)

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.

N o t e   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 1 is configured as shown below and you want to remove 0c0090-123456 from the Authorized Address list:

```
HP Switch(config)# show port-security 1
 Port Security

  Port : 1
  Learn Mode [Continuous] : Static        Address Limit [1] : 2
  Action [None] : None

  Authorized Addresses
  --------------------
  0c0090-123456
  0c0090-456456
```

When removing 0c0090-123456, first reduce the Address Limit by 1 to prevent the port from automatically adding another device that it detects on the network.
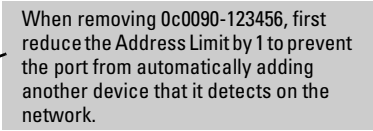
Figure 11-8. Example of Two Authorized Addresses on Port 1

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

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

The above command sequence results in the following configuration for port 1:

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

Figure 11-9. Example of Port 1 After Removing One MAC Address

MAC Lockdown

MAC Lockdown, also known as "static addressing," is the permanent assign-ment of a given MAC address (and VLAN, or Virtual Local Area Network) to a specific port on the switch. MAC Lockdown is used to prevent station movement and MAC address hijacking. It also controls address learning on the switch. When configured, the MAC Address can only be used on the assigned port and the client device will only be allowed on the assigned VLAN.

N o t e

Port security and MAC Lockdown are mutually exclusive on a given port. You can either use port security *or* MAC Lockdown, but never both at the same time on the same port.

Syntax: [no] static-mac < *mac-addr* > vlan < *vid* > interface < *port-number* >

You will need to enter a separate command for each MAC/VLAN pair you wish to lock down. If you do not specify a VLAN ID (VID) the switch inserts a VID of "1".

How It Works. When a device's MAC address is locked down to a port (typically in a pair with a VLAN) all information sent to that MAC address must go through the locked-down port. If the device is moved to another port it cannot receive data. Traffic to the designated MAC address goes only to the allowed port, whether the device is connected to it or not.

MAC Lockdown is useful for preventing an intruder from "hijacking" a MAC address from a known user in order to steal data. Without MAC Lockdown, this will cause the switch to learn the address on the malicious user's port, allowing the intruder to steal the traffic meant for the legitimate user.

MAC Lockdown ensures that traffic intended for a specific MAC address can only go through the one port which is supposed to be connected to that MAC address. It does not prevent intruders from transmitting packets with the locked MAC address, but it does prevent responses to those packets from going anywhere other than the locked-down port. Thus TCP connections cannot be established. Traffic sent to the locked address cannot be hijacked and directed out the port of the intruder.

If the device (computer, PDA, wireless device) is moved to a different port on the switch (by reconnecting the Ethernet cable or by moving the device to an area using a wireless access point connected to a different port on that same switch), the port will detect that the MAC Address is not on the appropriate port and will continue to send traffic out the port to which the address was locked.

Once a MAC address is configured for one port, you cannot perform port security using the same MAC address on any other port on that same switch.

You cannot lock down a single MAC Address/VLAN pair to more than one port; however you can lock down multiple different MAC Addresses to a single port on the same switch.

Stations can move from the port to which their MAC address is locked to other parts of the network. They can send, but will not receive data if that data must go through the locked down switch. Please note that if the device moves to a distant part of the network where data sent to its MAC address never goes through the locked down switch, it may be possible for the device to have full two-way communication. For full and complete lockdown network-wide all switches must be configured appropriately.

Other Useful Information. Once you lock down a MAC address/VLAN pair on one port that pair cannot be locked down on a different port.

You cannot perform MAC Lockdown and 802.1X authentication on the same port or on the same MAC address. MAC Lockdown and 802.1X authentication are mutually exclusive.

Lockdown is permitted on static trunks (manually configured link aggregations).

Differences Between MAC Lockdown and Port Security

Because port-security relies upon MAC addresses, it is often confused with the MAC Lockdown feature. However, MAC Lockdown is a completely different feature and is implemented on a different architecture level.

Port security maintains a list of allowed MAC addresses on a per-port basis. An address can exist on multiple ports of a switch. Port security deals with MAC addresses only while MAC Lockdown specifies both a MAC address and a VLAN for lockdown.

MAC Lockdown, on the other hand, is not a "list." It is a global parameter on the switch that takes precedence over any other security mechanism. The MAC Address will only be allowed to communicate using one specific port on the switch.

MAC Lockdown is a good replacement for port security to create tighter control over MAC addresses and which ports they are allowed to use (only one port per MAC Address on the same switch in the case of MAC Lockdown). (You can still use the port for other MAC addresses, but you cannot use the locked down MAC address on other ports.)

Using only port security the MAC Address could still be used on another port on the same switch. MAC Lockdown, on the other hand, is a clear one-to-one relationship between the MAC Address and the port. Once a MAC address has been locked down to a port it cannot be used on another port on the same switch.

The switch does not allow MAC Lockdown and port security on the same port.

MAC Lockdown Operating Notes

Limits. There is a limit of 500 MAC Lockdowns that you can safely code per switch. To truly lock down a MAC address it would be necessary to use the MAC Lockdown command for every MAC Address and VLAN ID on every switch. In reality few network administrators will go to this length, but it is important to note that just because you have locked down the MAC address and VID for a single switch, the device (or a hacker "spoofing" the MAC address for the device) may still be able to use another switch which hasn't been locked down.

Event Log Messages. If someone using a locked down MAC address is attempting to communicate using the wrong port the "move attempt" generates messages in the log file like this:

```
Move attempt (lockdown) logging:

W 10/30/03 21:33:43 maclock: Move 0001e6-1f96c0 to 15
denied

W 10/30/03 21:33:48 maclock: Move 0001e6-1f96c0 to 15
denied

W 10/30/03 21:33:48 maclock: Ceasing move-denied logs for
5m
```

These messages in the log file can be useful for troubleshooting problems. If you are trying to connect a device which has been locked down to the wrong port, it will not work but it will generate error messages like this to help you determine the problem.

Limiting the Frequency of Log Messages. The first move attempt (or intrusion) is logged as you see in the example above. Subsequent move attempts send a message to the log file. What this means is that the logging system checks again after the first 5 minutes to see if another attempt has been made to move to the wrong port. If this is the case the log file registers the most recent attempt and then checks again after one hour. If there are no further attempts in that period then it will continue to check every 5 minutes. If another attempt was made during the one hour period then the log resets itself to check once a day. The purpose of rate-limiting the log messaging is to prevent the log file from becoming too full. You can also configure the switch to send the same messages to a Syslog server. Refer to "Debug and Syslog Messaging Operation" in appendix C of the *Management and Configuration Guide* for your switch.

Deploying MAC Lockdown

When you deploy MAC Lockdown you need to consider how you use it within your network topology to ensure security. In some cases where you are using techniques such as "meshing" or Spanning Tree Protocol (STP) to speed up network performance by providing multiple paths for devices, using MAC Lockdown either will not work or else it defeats the purpose of having multiple data paths.

The purpose of using MAC Lockdown is to prevent a malicious user from "hijacking" an approved MAC address so they can steal data traffic being sent to that address.

As we have seen, MAC Lockdown can help prevent this type of hijacking by making sure that all traffic to a specific MAC address goes only to the proper port on a switch which is supposed to be connected to the real device bearing that MAC address.

However, you can run into trouble if you incorrectly try to deploy MAC Lockdown in a network that uses multiple path technology, like Spanning Tree or "mesh networks."

Let's examine a good use of MAC Lockdown within a network to ensure security first.
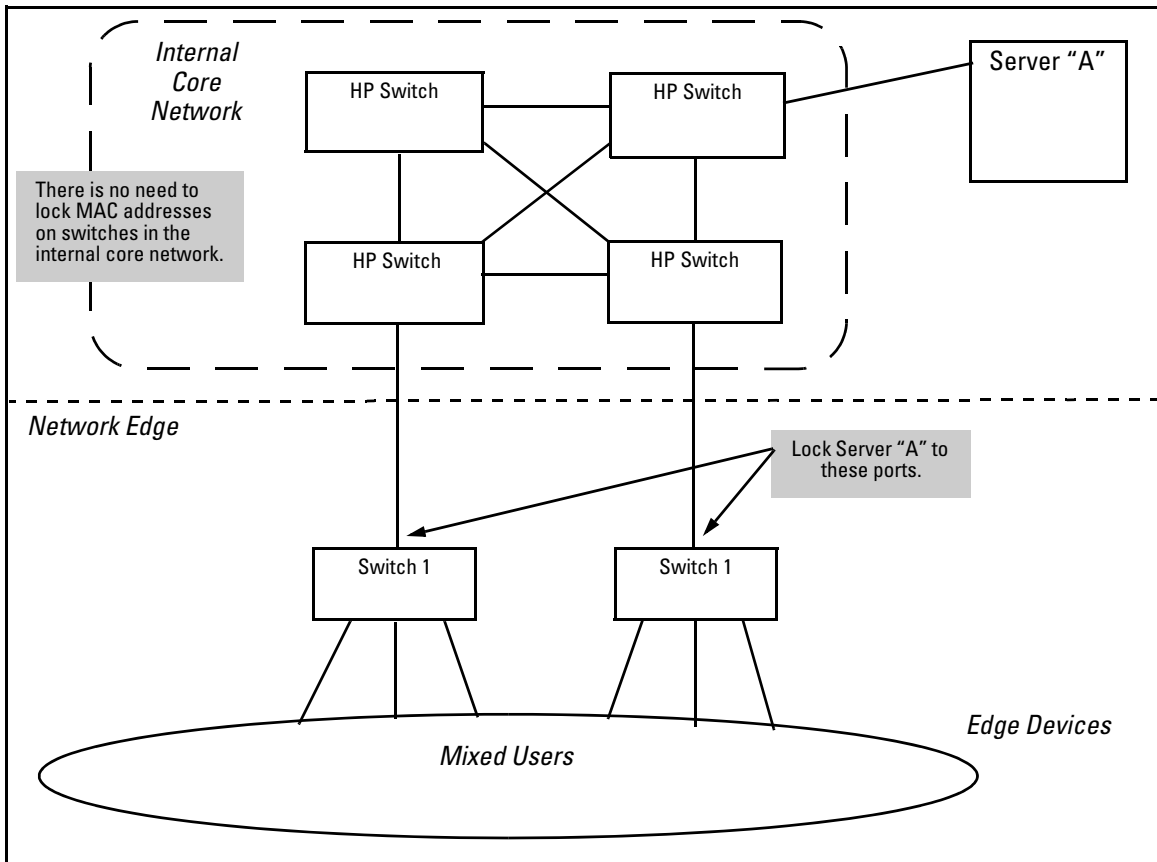
Figure 11-10. MAC Lockdown Deployed At the Network Edge Provides Security

Basic MAC Lockdown Deployment. In the Model Network Topology shown above, the switches that are connected to the edge of the network each have one and only one connection to the core network. This means each switch has only one path by which data can travel to Server A. You can use MAC Lockdown to specify that all traffic intended for Server A's MAC Address must go through the one port on the edge switches. That way, users on the edge can still use other network resources, but they cannot "spoof" Server A and hijack data traffic which is intended for that server alone.

The key points for this Model Topology are:

- The Core Network is separated from the edge by the use of switches which have been "locked down" for security.

- All switches connected to the edge (outside users) each have only one port they can use to connect to the Core Network and then to Server A.

- Each switch has been configured with MAC Lockdown so that the MAC Address for Server A has been locked down to one port per switch that can connect to the Core and Server A.

Using this setup Server A can be moved around within the core network, and yet MAC Lockdown will still prevent a user at the edge from hijacking its address and stealing data.

Please note that in this scenario a user with bad intentions at the edge can still "spoof" the address for Server A and send out data packets that look as though they came from Server A. The good news is that because MAC Lockdown has been used on the switches on the edge, any traffic that is sent *back* to Server A will be sent to the proper MAC Address because MAC Lockdown has been used. The switches at the edge will not send Server A's data packets anywhere but the port connected to Server A. (Data would not be allowed to go beyond the edge switches.)

Caution   Using MAC Lockdown still does not protect against a hijacker *within the core*! In order to protect against someone spoofing the MAC Address for Server A inside the Core Network, you would have to lock down each and every switch inside the Core Network as well, not just on the edge.

Problems Using MAC Lockdown in Networks With Multiple Paths. Now let's take a look at a network topology in which the use of MAC Lockdown presents a problem. In the next figure, Switch 1 (on the bottom-left) is located at the edge of the network where there is a mixed audience that might contain hackers or other malicious users. Switch 1 has two paths it could use to connect to Server A. If you try to use MAC Lockdown here to make sure that all data to Server A is "locked down" to one path, connectivity problems would be the result since both paths need to be usable in case one of them fails.
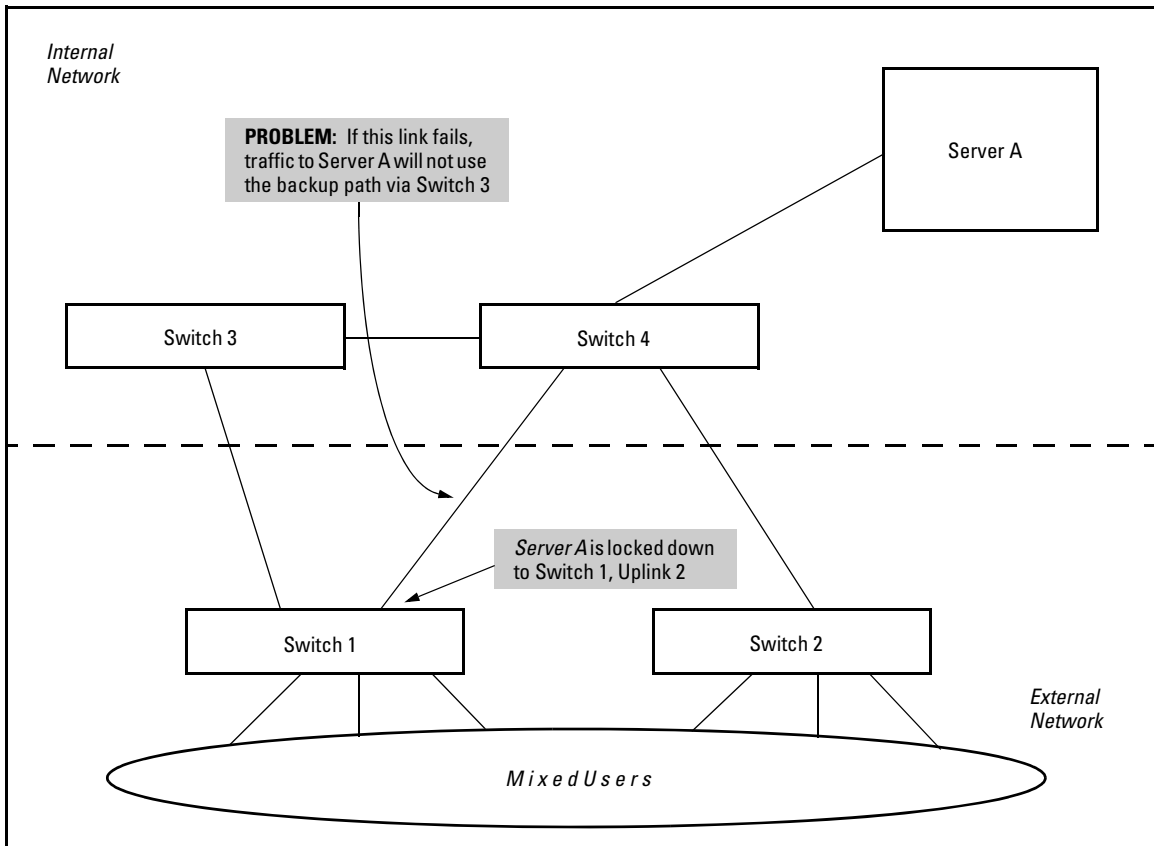
Figure 11-11. Connectivity Problems Using MAC Lockdown with Multiple Paths

The resultant connectivity issues would prevent you from locking down Server A to Switch 1. And when you remove the MAC Lockdown from Switch 1 (to prevent broadcast storms or other connectivity issues), you then open the network to security problems. The use of MAC Lockdown as shown in the above figure would defeat the purpose of using MSTP or having an alternate path.

Technologies such as MSTP or "meshing" are primarily intended for an internal campus network environment in which all users are trusted. MSTP and "meshing" do not work well with MAC Lockdown.

If you deploy MAC Lockdown as shown in the Model Topology in figure 11-10 (page 11-23), you should have no problems with either security or connectivity.

MAC Lockout

MAC Lockout involves configuring a MAC address on all ports and VLANs for a switch so that any traffic to or from the "locked-out" MAC address will be dropped. This means that all data packets addressed to or from the given address are stopped by the switch. MAC Lockout is implemented on a per switch assignment.

You can think of MAC Lockout as a simple blacklist. The MAC address is locked out on the switch and on all VLANs. No data goes out or in from the blacklisted MAC address to a switch using MAC Lockout.

To fully lock out a MAC address from the network it would be necessary to use the MAC Lockout command on all switches.

To use MAC Lockout you must first know the MAC Address you wish to block.

Syntax: [no] lockout-mac < *mac-address* >

> Lockout the specified MAC address.

How It Works. Let's say a customer knows there are unauthorized wireless clients who should not have access to the network. The network administrator "locks out" the MAC addresses for the wireless clients by using the MAC Lockout command (**lockout-mac <*mac-address*>**). When the wireless clients then attempt to use the network, the switch recognizes the intruding MAC addresses and prevents them from sending or receiving data on that network.

If a particular MAC address can be identified as unwanted on the switch then that MAC Address can be disallowed on all ports on that switch with a single command. You don't have to configure every single port—just perform the command on the switch and it is effective for all ports.

MAC Lockout overrides MAC Lockdown, port security, and 802.1X authentication.

You cannot use MAC Lockout to lock:

- Broadcast or Multicast Addresses (Switches do not learn these)
- Switch Agents (The switch's own MAC Address)

There are limits for the number of VLANs, Multicast Filters, and Lockout MACs that can be configured concurrently as all use MAC table entries. The limits are shown below.

Table 11-1.   Limits on Lockout MACs

# VLANs	# Multicast Filters	# Lockout MACs
< = 1024	16	16
1025-2048	8	8

If someone using a locked out MAC address tries to send data through the switch a message is generated in the log file:

```
Lockout logging format:

W 10/30/03 21:35:15 maclock: 0001e6-1f96c0 detected on
port 15
W 10/30/03 21:35:18 maclock: 0001e6-1f96c0 detected on
port 15
W 10/30/03 21:35:18 maclock: Ceasing lock-out logs for
5m
```

As with MAC Lockdown a rate limiting algorithm is used on the log file so that it does not become overclogged with error messages. (Refer to "Limiting the Frequency of Log Messages" on page 11-22.)

Port Security and MAC Lockout

MAC Lockout is independent of port-security and in fact will override it. MAC Lockout is preferable to port-security to stop access from known devices because it can be configured for all ports on the switch with one command.

It is possible to use MAC Lockout in conjunction with port-security. You can use MAC Lockout to lock out a single address—deny access to a specific device—but still allow the switch some flexibility in learning other MAC Addresses. Be careful if you use both together, however:

- If a MAC Address is locked out and appears in a static learn table in port-security, the apparently "authorized" address will still be locked out anyway.

- MAC entry configurations set by port security will be kept even if MAC Lockout is configured and the original port security settings will be honored once the Lockout is removed.

- A port security static address is permitted to be a lockout address. In that case (MAC Lockout), the address will be locked out (SA/DA drop) even though it's an "authorized" address from the perspective of port security.

- When MAC Lockout entries are deleted, port security will then re-learn the address as needed later on.

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.

Denial of Service Packet Filtering

When enabled, Denial of Service packet filtering automatically scans traffic passing through the switch for well-known frames that are used to conduct denial-of-service attacks on network devices. When a packet is identified as a threat, it is automatically dropped.

Syntax: [no] dos-filter

> *Enables traffic scanning for well known frames used to conduct denial of service (dos) attacks*

> *The* **no** *form of the command disables traffic scanning.*

To display the status of DOS filtering, use the **show** command.

```
HP Switch(config)# show dos-filter

  Denial of Service packet filtering: Enabled
```

Figure 11-12. Example of the show dos-filter Command Output

Web Browser Support

The Web browser interface has an on/off setting to enable or disable DOS packet filtering. Click on the **Security** tab and select the **DOS Filter** button.

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 port-security 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 network management applications such as IMC via an SNMP trap sent to a network 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 1 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 1.

```
Status and Counters - Intrusion Log
 Port  MAC Address         Date / Time
 ----  ------------  -------------------------
 1     080009-e93d4f         03/07/06 21:09:34
 1     080009-e93d4f         03/07/06 10:18:43
```

Figure 11-13. 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.

Note on Send-Disable Operation	On a given port, if the intrusion action is to send an SNMP trap and then disable the port (**send-disable**), and an intruder is detected on the port, then the switch sends an SNMP trap, sets the port's alert flag, and disables the port. If you re-enable the port without resetting the port's alert flag, then the port operates as follows:

■ The port comes up and will block traffic from unauthorized devices it detects.

■ If the port detects another intruder, it will send another SNMP trap, but will not become disabled again unless you first reset the port's intrusion flag.

This operation enables the port to continue passing traffic for authorized devices while you take the time to locate and eliminate the intruder. Otherwise, the presence of an intruder could cause the switch to repeatedly disable the 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
   4. Port Status

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

                        Intrusion                              Flow
          Port    Type    Alert    Enabled   Status    Mode     Ctrl
          ----  --------  -------  -------   ------   --------  ---------
           1    10/100TX   No        Yes      Up       Auto     off
           2    10/100TX   No        Yes      Up       Auto     off
           3    10/100TX   Yes       Yes      Up       Auto     off
           4    10/100TX   No        Yes      Up       Auto     off
           5    10/100TX   No        Yes      Up       Auto     off
           6    10/100TX   No        Yes      Down     Auto     off
           7    10/100TX   No        Yes      Up       Auto     off
           8    10/100TX   No        Yes      Down     Auto     off


          Actions->   Back       Intrusion log     Help

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

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

Figure 11-14. Example of Port Status Screen with Intrusion Alert on Port 3

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

```
=========================- CONSOLE - MANAGER MODE -=========================
                        Status and Counters - Intrusion Log

          Port   MAC Address          Date / Time
          ----  -------------      ---------------------
           3    080009-6563e2             08/08/02 16:58:02
           1    0060b0-896e00             08/08/02 15:28:21
           3    080009-cf558f   prior to 08/08/02 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.
```

MAC Address of Intruding Device on Port 3

System Time of Intrusion on Port 3

Indicates this intrusion on port 3 occurred prior to a reset (reboot) at the indicated time and date.

Figure 11-15. Example of the Intrusion Log Display

The example in Figure 7-11 shows two intrusions for port 3 and one intrusion for port 1. In this case, only the most recent intrusion at port 3 has not been acknowledged (reset). This is indicated by the following:

• Because the Port Status screen (figure 11-14 on page 11-32) does not indicate an intrusion for port 1, the alert flag for the intrusion on port 1 has already been reset.

• Since the switch can show only one uncleared intrusion per port, the alert flag for the older intrusion for port 3 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 3 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 3 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 3 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 11-15, 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, refer to "Operating Notes for Port Security" on page 11-36.)

Syntax: show interfaces brief

List intrusion alert status (and other port status informa-tion)'.

show port-security intrusion-log

List intrusion log content.

clear intrusion-flags

Clear intrusion flags on all ports.

port-security [e] < port-number > clear-intrusion-flag

Clear the intrusion flag on one or more specific ports.

In the following example, executing **show interfaces brief** lists the switch's port status, which indicates an intrusion alert on port 1.

```
   HP Switch# show interfaces brief        Intrusion Alert on port 1.
    Status and Counters - Port Status
                      | Intrusion                            Flow
     Port   Type      |  Alert    Enabled  Status   Mode     Ctrl
     ----  --------   + --------- -------  ------  ---------- ---------
      1    10/100TX   |  Yes   ◄── Yes      Up     10HDx      off
      2    10/100TX   |  No        Yes      Up     10HDx      off
      3    10/100TX   |  No        Yes      Up     10HDx      off
      4    10/100TX   |  No        Yes      Up     10HDx      off
```

Figure 11-16.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 port-security intrusion-log** command. For example:

MAC Address of latest Intruder on Port 1

Earlier intrusions on port 1 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

```
   HP Switch# show port-security intrusion-log

   Status and Counters - Intrusion Log

   Port   MAC Address             Date / Time
   ----  ------------  -------------------------
    1     080009-e93d4f           07/03/06 21:09:34
    1     080009-21ae84           07/03/06 17:26:27
    1     080009-e93d4f prior to 07/03/06 17:18:43
```

Figure 11-17. Example of the Intrusion Log with Multiple Entries for the Same Port

The above example shows three intrusions for port 1. 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 < *port-list* > 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 1 and enable the switch to enter any subsequent intrusion for port 1 in the Intrusion Log, execute the port-security **clear-intrusion-flag** command. If you then re-display the port status screen, you will see that the Intrusion Alert entry for port 1 has changed to "**No**". (Executing **show port-security intrusion-log** again will result in the same display as above, and does not include the Intrusion Alert status.)

```
HP Switch(config)# port-security 1 clear-intrusion-flag
HP Switch(config)# show interfaces brief
```

```
 Status and Counters - Port Status          Intrusion Alert on port 1 is now cleared.
                     | Intrusion                              Flow   Bcast
   Port    Type      |   Alert     Enabled Status    Mode     Ctrl   Limit
   ----  --------  + ---------  ------- ------ ---------- --------- ------
    1    10/100TX  | No          Yes     Up     10HDx     off      0
    2    10/100TX  | No          Yes     Up     10HDx     off      0
    3    10/100TX  | No          Yes     Up     10HDx     off      0
```

Figure 11-18. Example of Port Status Screen After Alert Flags Reset

For more on clearing intrusions, see "Note on Send-Disable Operation" on page 11-31

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 3 — 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, display the Intrusion Log, as shown below.

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

Log Listing with No Security Violation Detected

Log Command with "security" for Search String

```
HP Switch(config)# log security
Keys: W=Uarning    I=Information
        M=Major       D=Debug
----   Event Log listing: Events Since Boot   ----
W 08/01/02 01:18:15 FFI: port 2 - Security Violation
W 08/01/02 04:28:03 FFI: port 1 - Security Violation
---- Bottom of Log : Events Listed = 2 ----

HP Switch(config)# log security
Keys: W=Uarning    I=Information
        M=Major       D=Debug
----   Event Log listing: Events Since Boot   ----
----   Bottom of Log : Events Listed = 0   ----
```

Figure 11-19. Example of Log Listing With and Without Detected Security Violations

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 **[Overview]** button. If there is a "Security Violation" entry, do the following:

   a. Click on the **Security** tab.

   b. Click on **[Intrusion Log]**. "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 **[Reset Alert Flags]**.

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 IMC to manage your network, you can use the device properties page to link MAC addresses to their corresponding IP addresses.

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.

LACP Not Available on Ports Configured for Port Security. To maintain security, LACP is not allowed on ports configured for port security. If you configure port security on a port on which LACP (active or passive) is configured, the switch removes the LACP configuration, displays a notice that LACP is disabled on the port(s), and enables port security on that port. For example:

```
HP Switch(config)# port-security e 17 learn-mode static address-
limit 2
LACP has been disabled on secured port(s).
HP Switch(config)#
```

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

```
HP Switch(config)# int e 17 lacp passive
Error configuring port 17: LACP and port security cannot be run
together.
HP Switch(config)#
```

To restore LACP to the port, you must remove port security and re-enable LACP active or passive.

Configuring Protected Ports

There are situations where you want to provide internet access to users but prevent them from accessing each other. To achieve this control, you can use the **protected-ports** command. The command applies per-port, and filters the outbound traffic from a port. This allows the configuration of two port groups on a switch—protected ports and unprotected ports. The ports have these characteristics:

■ Traffic from protected ports is not forwarded to other protected ports.

■ Protected ports can communicate with unprotected ports, but not with each other.

■ Unprotected ports can communicate with all ports.

■ The **protected-ports** command applies to logical ports (trunks as well as untrunked ports)

Syntax: [no] protected-ports <port-list>

Prevents the selected ports from communicating with each other.
Default: All ports unprotected.

no protected-ports all

Clears the protection from all ports; all ports can now communicate with each other.

```
HP Switch(config)# protected-ports 4-5
```

Figure 11-20. Example of Protected Ports Command for Ports 4 and 5

To display information about which ports have been configured as protected ports, enter this command:

```
HP Switch(config)# show protected-ports
```

```
HP Switch(config)# show protected-ports

  Protected-ports  : 4-5
  Unprotected-ports: 1-3,6-26
```

Figure 11-21. Example Showing Protected Ports and Unprotected Ports

If you display the running config file (**show running-config**) you will see the ports that have been selected as protected ports.

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

Running configuration:

; J9229A Configuration Editor; Created on release #YA.15.XX

hostname "Switch"
snmp-server community "public" Unrestricted
vlan 1
    name "DEFAULT_VLAN"
    untagged 1-26
    ip address dhcp-bootp
    exit
protected-ports 4-5          ◄——— Protected ports
```

Figure 11-22. Example of Running Config File Showing Protected Ports

In the example in Figure 11-23, ports 1 through 8 are protected. Port 2 connects to Room 2 and Port 4 connects to Room 4. Ports 9 and 10 are unprotected and provide access to the internet. Ports 2 and 4 are able to access the internet, but are not able to communicate with each other or any of the other rooms that are connected to protected ports.
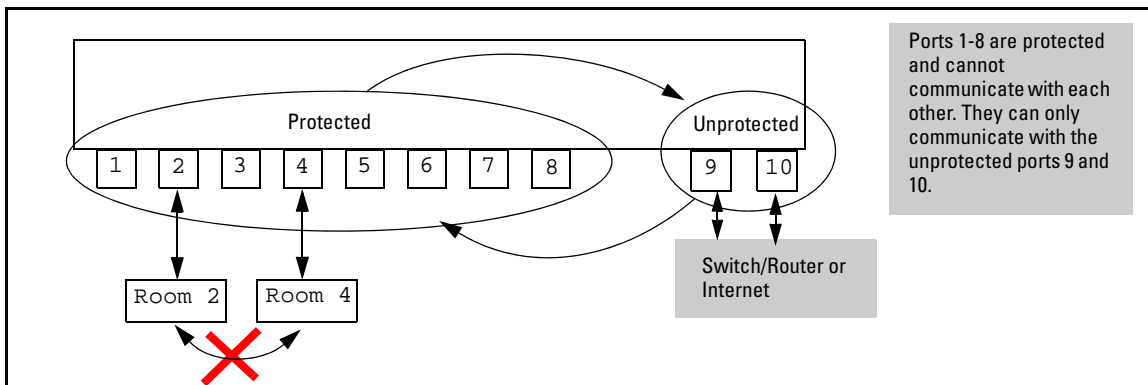
Figure 11-23. Example With Ports 1-8 Protected and Ports 9 and 10 Unprotected

Using Authorized IP Managers

Overview

Authorized IP Manager Features

Feature	Default	Menu	CLI	Web
Listing (Showing) Authorized Managers	n/a	page 12-4	page 12-5	page 12-9
Configuring Authorized IP Managers	None	page 12-4	page 12-5	page 12-9
Building IP Masks	n/a	page 12-11	page 12-11	page 12-11
Operating and Troubleshooting Notes	n/a	page 12-14	page 12-14	page 12-14

The Authorized IP Managers feature uses IP addresses and masks to deter-mine which stations (PCs or workstations) can access the switch through the network. This covers access through the following means:

– Telnet and other terminal emulation applications
– The switch's web browser interface
– SSH
– SNMP versions 1, 2 and 3(with a correct community name)
– TFTP

Also, when configured in the switch, the Authorized IP Managers feature takes precedence over local passwords, TACACS+, RADIUS, Port-Based Access Control (802.1X), and Port Security. This means that the IP address of a networked management device must be authorized before the switch will attempt to authenticate the device by invoking any other access security features. If the Authorized IP Managers feature disallows access to the device, then access is denied. Thus, with authorized IP managers configured, having the correct passwords is not sufficient for accessing the switch through the network unless the station attempting access is also included in the switch's Authorized IP Managers configuration.

You can use Authorized IP Managers along with other access security features to provide a more comprehensive security fabric than if you use only one or two security options.

Note

When no Authorized IP manager rules are configured, the access method feature is disabled, that is, access is not denied.

Configuration Options

You can configure:

■ Up to 10 authorized manager *addresses*, where each address applies
to either a single management station or a group of stations

■ Manager or Operator access privileges

Caution Configuring Authorized IP Managers does not protect access to the switch
through a modem or direct connection to the Console (RS-232) port. Also, if
an unauthorized station "spoofs" an authorized IP address, it can gain man-
agement access to the switch even though a duplicate IP address condition
exists. For these reasons, you should enhance your network's security by
keeping physical access to the switch restricted to authorized personnel, using
the username/password and other security features available in the switch,
and preventing unauthorized access to data on your management stations.

Access Levels

For each authorized manager address, you can configure either of these access
levels:

■ **Manager:** Enables full access to all screens for viewing, configuration,
and all other operations available.

■ **Operator:** Allows read-only access. (This is the same access that is
allowed by the switch's operator-level password feature.)

Defining Authorized Management
Stations

■ **Authorizing Single Stations:** The table entry authorizes a single man-agement station to have IP access to the switch. To use this method, just enter the IP address of an authorized management station in the Autho-rized Manager IP column, and leave the IP Mask set to **255.255.255.255**. This is the easiest way to use the Authorized Managers feature. (For more on this topic, see "Configuring One Station Per Authorized Manager IP Entry" on page 12-11.)

■ **Authorizing Multiple Stations:** The table entry uses the IP Mask to authorize access to the switch from a defined group of stations. This is useful if you want to easily authorize several stations to have access to the switch without having to type in an entry for every station. All stations in the group defined by the one Authorized Manager IP table entry and its associated IP mask will have the same access level—Manager or Operator. (For more on this topic, refer to "Configuring Multiple Stations Per Authorized Manager IP Entry" on page 12-11.)

To configure the switch for authorized manager access, enter the appropriate *Authorized Manager IP* value, specify an *IP Mask*, and select either **Manager** or **Operator** for the *Access Level*. The IP Mask determines how the Authorized Manager IP value is used to allow or deny access to the switch by a manage-ment station.

Note   If the management VLAN is configured, access can only be on that VLAN.

Overview of IP Mask Operation

The default IP Mask is 255.255.255.255 and allows switch access only to a station having an IP address that is identical to the Authorized Manager IP parameter value. ("255" in an octet of the mask means that only the exact value in the corresponding octet of the Authorized Manager IP parameter is allowed in the IP address of an authorized management station.) However, you can alter the mask and the Authorized Manager IP parameter to specify ranges of authorized IP addresses. For example, a mask of **255.255.255.0** and any value for the Authorized Manager IP parameter allows a range of 0 through 255 in the 4th octet of the authorized IP address, which enables a block of up to 254 IP addresses for IP management access (excluding 0 for the network and 255 for broadcasts). A mask of **255.255.255.252** uses the 4th octet of a given Autho-rized Manager IP address to authorize four IP addresses for management station access. The details on how to use IP masks are provided under "Building IP Masks" on page 12-11.

N o t e	The IP Mask is a method for recognizing whether a given IP address is authorized for management access to the switch. This mask serves a different purpose than IP subnet masks and is applied in a different manner.

Menu: Viewing and Configuring IP Authorized Managers

Only IPv4 is supported when using the menu to set the management access method.

From the console Main Menu, select:

2. Switch Configuration …

    6. IP Authorized Managers

```
==========================- CONSOLE - MANAGER MODE -============================
                    Switch Configuration - IP Managers

  Authorized Manager IP         IP Mask           Access Level  Access Method
  ---------------------   ---------------------   ------------  -------------
  10.10.240.2             255.255.255.255         Manager       all
  10.10.245.3             255.255.255.255         Operator      ssh
  10.10.246.200           255.255.255.255         Operator      tftp
  10.10.245.30            255.255.255.0           Operator      ssh
                                            1. Select Add to add an authorized manager
                                               to the list.


  Actions->   Back     Add      Edit      Delete     Help
```

Figure 12-1. Example of How to Add an Authorized Manager Entry

```
==========================- CONSOLE - MANAGER MODE -============================
                    Switch Configuration - IP Managers
                                            Enter an Authorized Manager IP address here.
  Authorized Manager IP: 10.10.245.3
  IP Mask [255.255.255.255]:255.255.255.255   Use the default mask to allow access by one
  Access Level:Operator                       management device, or edit the mask to allow
                                              access by a block of management devices. See
                                              "Building IP Masks" on page 12-11.

                                              Use the Space bar to select Manager or Operator
                                              access.

                                              Press [Enter], then [S] (for Save) to configure the
                                              IP Authorized Manager entry.
  Actions->   Back     Add      Edit      Delete     Help
```

Figure 12-2. Example of Edit Menu for Authorized IP Managers

Editing or Deleting an Authorized Manager Entry. Go to the IP Managers List screen (figure 12-1), highlight the desired entry, and press **[E]** (for **Edit**) or **[D]** (for **Delete**).

CLI: Viewing and Configuring Authorized IP Managers

Authorized IP Managers Commands Used in This Section

Command	Page
show ip authorized-managers	below
ip authorized-managers	12-6
<ip-address>	12-6
mask *<mask-bits>*	12-6
access <operator	manager>	12-6
access-method	12-5

Listing the Switch's Current Authorized IP Manager(s)

Use the **show ip authorized-managers** command to list IP stations authorized to access the switch. For example:

```
HP Switch(config)# show ip authorized-manager

IPV4 Authorized Managers
-----------------------

 Address : 10.10.10.10
 Mask    : 255.255.255.255
 Access  : Manager
```

Figure 12-3. Example of show authorized-managers Command

Configuring IP Authorized Managers for the Switch

See the *IPv6 Configuration Guide* for information about Authorized IP
manager configuration with IPv6 addresses.

Syntax: [no] ip authorized-managers <*ip-address*> <*ip-mask*>> access [manager
| operator] access-method [all | ssh | telnet | web | snmp | tftp]

> *Configures one or more authorized IP addresses.*

access [manager | operator]

> *Configures the privilege level for* <ip-address>. *Applies
> only to access through telnet, SSH, SNMPv1, SNMPv2c,
> and SNMPv3.*
> *Default:* **manager**

access-method [all | ssh | telnet | web | snmp | tftp]

> *Configures access levels by access method and IP
> address. Each management method can have its own set
> of authorized managers.*
> *Default:* **all**

```
HP Switch(config)# ip authorized-managers 10.10.10.2 255.255.255.255 manager
```

Figure 12-4. Example of Configuring IP Authorized Manager

To Authorize Manager Access. This command authorizes manager-level
access for any station with an IP address of 10.28.227.0 through 10.28.227.255:

```
HPswitch(config)# ip authorized-managers 10.28.227.101
255.255.255.0 access manager
```

Similarly, the next command authorizes manager-level access for any station
having an IP address of 10.28.227.101 through 103:

```
HPswitch(config)# ip authorized-managers 10.28.227.101
255.255.255.252 access manager
```

If you omit the <*mask bits*> when adding a new authorized manager, the switch
automatically uses **255.255.255.255**. If you do not specify either Manager or
Operator access, the switch assigns the Manager access.

To Edit an Existing Manager Access Entry. To change the mask or
access level for an existing entry, use the entry's IP address and enter the new
value(s). (Notice that any parameters not included in the command will be set
to their default.):

```
HPswitch(config)# ip authorized-managers
        10.28.227.101 255.255.255.0 access operator
```

The above command replaces the existing mask and access level for IP address 10.28.227.101 with `255.255.255.0` and `operator`.

The following command replaces the existing mask and access level for IP address 10.28.227.101 with `255.255.255.255` and `manager` (the defaults) because the command does not specify either of these parameters.

```
HPswitch(config)# ip authorized-managers 10.28.227.101
```

To Delete an Authorized Manager Entry. This command uses the IP address of the authorized manager you want to delete:

```
HPswitch(config)# no ip authorized-managers 10.28.227.101
```

Setting the Management Access Method

This feature allows the configuration of access methods for IP Authorized Manager entries. Each of the management access methods will have its own set of authorized managers. The access methods include:

- SSH
- Telnet
- Web
- TFTP
- SNMP

You can configure the access method via the CLI, the menu, or through the web interface. The menu interface only supports IPv4. The following restrictions apply to all three methods of configuration.

- When no IP authorized manager rules are configured, the access method feature is disabled, that is, access is not denied.

- If the Management VLAN is configured, access can only be on that VLAN.

- Using the access method feature is optional. If no access method is configured, the access method defaults to "all".

- If access is not specified, it defaults to "manager".

- The IP mask defaults to 255.255.255.255.

```
HP Switch(config)# ip authorized-managers 10.10.10.2 255.255.255.255 manager
                   access-method ssh
```

Figure 12-5. Example of Configuring IP Authorized Manager Access Method SSH

```
HP Switch(config)# show ip authorized-manager

IPV4 Authorized Managers
------------------------

 Address : 10.10.10.10
 Mask    : 255.255.255.255
 Access  : Manager
 Access Method : ssh
```

Figure 12-6. Example of show authorized-managers Command with Access Method Configured

Setting the Management Access Method—Menu

Only IPv4 is supported when using the menu to set the management access method.

To access the menu screen, type **menu** at the switch prompt, then select **2. Switch Configuration**, then **6. IP Authorized Managers**. The menu screen for IP Managers displays. Click on **Edit** to make changes.

```
==========================- CONSOLE - MANAGER MODE -===========================
                  Switch Configuration - IP Managers

  Authorized Manager IP          IP Mask            Access Level  Access Method
  --------------------      ----------------------  -----------   -------------
  10.10.240.2               255.255.255.255         Manager       all
  10.10.245.3               255.255.255.255         Operator      ssh
  10.10.246.200             255.255.255.255         Operator      tftp
  10.10.245.30              255.255.255.0           Operator      ssh



  Actions->    Back     Add     Edit     Delete     Help
```

Figure 12-7. Example of Menu Showing Authorized Managers with Access Method

```
==========================- CONSOLE - MANAGER MODE -===========================
                  Switch Configuration - IP Managers

  Authorized Manager IP: 10.10.245.3
  IP Mask [255.255.255.255]:255.255.255.255
  Access Level:Operator
  Access Method:ssh


  Actions->    Back     Add     Edit     Delete     Help
```

Figure 12-8. Example of Edit Menu for IP Managers

Setting the Management Access Method—Web Interface

To set the management access method in the web interface, click on the **Security** tab, and then click on the **Authorized Addresses** button. Fill in the fields with the correct information and click **Add**.

The Authorized Managers IP list in the web interface is the same list that was configured with the **ip authorized-managers** command in the CLI.
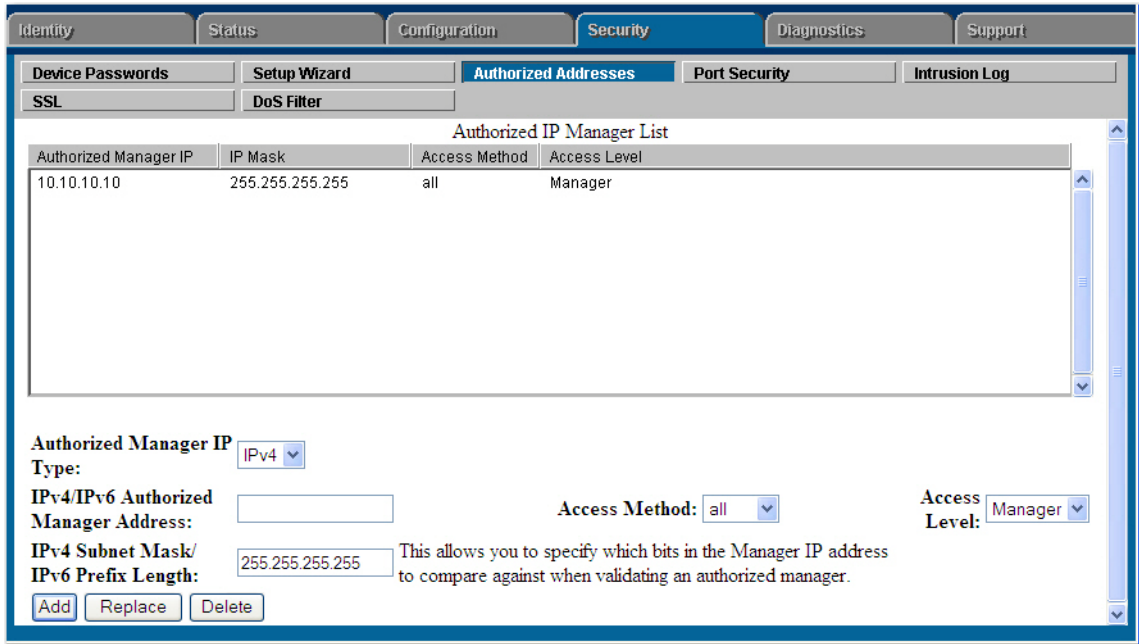
Figure 12-9. Example of Configuring Authorized Manager Access Method in the Web Interface

Web Proxy Servers

If you use the web browser interface to access the switch from an authorized IP manager station, it is highly recommended that you avoid using a web proxy server in the path between the station and the switch. This is because switch access through a web proxy server requires that you first add the web proxy server to the Authorized Manager IP list. This reduces security by opening switch access to anyone who uses the web proxy server.

How to Eliminate the Web Proxy Server

There are two ways to eliminate a web proxy server from the path between a station and the switch:

1. Add the IP address or DNS name of the switch to the non-proxy or Exceptions list in the web browser interface used on the authorized station.

2. If you don't need proxy server access on the authorized station, disable the proxy server feature in the station's web browser interface.

N o t e   IP or MAC authentication can be used without a web proxy server.

Using a Web Proxy Server to Access the Web Browser Interface

C a u t i o n   This is NOT recommended. Using a web proxy server between the stations and the switch poses a security risk. If the station uses a web proxy server to connect to the switch, any proxy user can access the switch.

If it is necessary to use the switch's web browser interface and your browser access is through a web proxy server, perform these steps:

1. Enter the web proxy server's MAC address in the port's Authorized Addresses list.

2. Enter the web proxy server's IP address in the switch's IP Authorized Managers list.

You must perform both of these steps or the switch only detects the proxy server's MAC address and IP address instead of your workstation addresses, and your connection is considered unauthorized.

Web-Based Help

For web-based help on how to use the web browser interface screen, click on the **[?]** button provided on the web browser screen.

Building IP Masks

The IP Mask parameter controls how the switch uses an Authorized Manager IP value to recognize the IP addresses of authorized manager stations on your network.

Configuring One Station Per Authorized Manager IP Entry

This is the easiest way to apply a mask. If you have ten or fewer management and/or operator stations, you can configure them by adding the address of each to the Authorized Manager IP list with **255.255.255.255** for the corresponding mask. For example, as shown in figure 12-3 on page 12-5, if you configure an IP address of **10.28.227.125** with an IP mask of **255.255.255.255**, only a station having an IP address of **10.28.227.125** has management access to the switch.

Table 9-1.   Analysis of IP Mask for Single-Station Entries

	1st Octet	2nd Octet	3rd Octet	4th Octet	**Manager-Level or Operator-Level Device Access**
IP Mask	255	255	255	255	The "255" in each octet of the mask specifies that only the exact value in that octet of the corresponding IP address is allowed. This mask allows management access only to a station having an IP address of 10.33.248.5.
Authorized Manager IP	10	33	248	5	

Configuring Multiple Stations Per Authorized Manager IP Entry

The mask determines whether the IP address of a station on the network meets the criteria you specify. That is, for a given Authorized Manager entry, the switch applies the IP mask to the IP address you specify to determine a range of authorized IP addresses for management access. As described above, that range can be as small as one IP address (if **255** is set for all octets in the mask), or can include multiple IP addresses (if one or more octets in the mask are set to less than **255**).

If a bit in an octet of the mask is "on" (set to 1), then the corresponding bit in the IP address of a potentially authorized station must match the same bit in the IP address you entered in the Authorized Manager IP list. Conversely, if a bit in an octet of the mask is "off" (set to 0), then the corresponding bit in the IP address of a potentially authorized station on the network does not have to match its counterpart in the IP address you entered in the Authorized Manager

IP list. Thus, in the example shown above, a "255" in an IP Mask octet (*all* bits in the octet are "on") means only one value is allowed for that octet—the value you specify in the corresponding octet of the Authorized Manager IP list. A "0" (all bits in the octet are "off") means that any value from 0 to 255 is allowed in the corresponding octet in the IP address of an authorized station. You can also specify a series of values that are a subset of the 0-255 range by using a value that is greater than 0, but less than 255.

Table 9-2.   Analysis of IP Mask for Multiple-Station Entries

	1st Octet	2nd Octet	3rd Octet	4th Octet	Manager-Level or Operator-Level Device Access
IP Mask	255	255	255	0	The "255" in the first three octets of the mask specify that only the exact value in the octet of the corresponding IP address is allowed. However, the zero (0) in the 4th octet of the mask allows any value between 0 and 255 in that octet of the corresponding IP address. This mask allows switch access to any device having an IP address of 10.28.227.*xxx*, where *xxx* is any value from 0 to 255.
Authorized Manager IP	10	28	227	125	
IP Mask	255	255	255	249	In this example (figure 9-3, below), the IP mask allows a group of up to 4 management stations to access the switch. This is useful if the only devices in the IP address group allowed by the mask are management stations. The "249" in the 4th octet means that bits 0 and 3 - 7 of the 4th octet are fixed. Conversely, bits 1 and 2 of the 4th octet are variable. Any value that matches the authorized IP address settings for the fixed bits is allowed for the purposes of IP management station access to the switch. Thus, any management station having an IP address of 10.28.227.<u>121</u>, <u>123</u>, <u>125</u>, or <u>127</u> can access the switch.
Authorized IP Address	10	28	227	125	

Table 9-3. Example of How the Bitmap in the IP Mask Defines Authorized Manager Addresses

4th Octet of IP Mask:				249				
4th Octet of Authorized IP Address:				**5**				

Bit Numbers	Bit 7	Bit 6	Bit 5	Bit 4	Bit 3	Bit 2	Bit 1	Bit 0	
Bit Values	128	64	32	16	8	4	2	1	
4th Octet of IP Mask (249)	■	■	■	■	■	□	□	■	Bits 1 and 2 in the mask are "off", and bits 0 and 3 - 7 are "on", creating a value of 249 in the 4th octet.
4th Octet of IP Authorized Address (125)	□	■	■	■	■	■	□	■	Where a mask bit is "on", the corresponding bit setting in the address of a potentially authorized station must match the IP Authorized Address setting for that same bit. Where a mask bit is "off" the corresponding bit setting in the address can be either "on" or "off". In this example, in order for a station to be authorized to access the switch:

- The first three octets of the station's IP address must match the Authorized IP Address.
- Bit 0 and Bits 3 through 6 of the 4th octet in the station's address must be "on" (value = 1).
- Bit 7 of the 4th octet in the station's address must be "off" (value = 0).
- Bits 1 and 2 can be either "on" or "off".

This means that stations with the IP address 13.28.227.*X* (where *X* is 121, 123, 125, or 127) are authorized.

Additional Examples for Authorizing Multiple Stations

	Entries for Authorized Manager List				Results
IP Mask	255	255	0	255	This combination specifies an authorized IP address of 10.33.*xxx*.1. It could be applied, for example, to a subnetted network where each subnet is defined by the third octet and includes a management station defined by the value of "1" in the fourth octet of the station's IP address.
Authorized Manager IP	10	33	248	1	
IP Mask	255	238	255	250	Allows 230, 231, 246, and 247 in the 2nd octet, and 194, 195, 198, 199 in the 4th octet.
Authorized Manager IP	10	247	100	195	

Operating Notes

■ **Network Security Precautions:** You can enhance your network's security by keeping physical access to the switch restricted to authorized personnel, using the password features built into the switch, using the additional security features described in this manual, and preventing unauthorized access to data on your management stations.

■ **Modem and Direct Console Access:** Configuring authorized IP managers does not protect against access to the switch through a modem or direct Console (RS-232) port connection.

■ **Duplicate IP Addresses:** If the IP address configured in an authorized management station is also configured (or "spoofed") in another station, the other station can gain management access to the switch even though a duplicate IP address condition exists.

■ **Web Proxy Servers:** If you use the web browser interface to access the switch from an authorized IP manager station, it is recommended that you avoid the use of a web proxy server in the path between the station and the switch. This is because switch access through a web proxy server requires that you first add the web proxy server to the Authorized Manager IP list. *This reduces security by opening switch access to anyone who uses the web proxy server.* The following two options outline how to eliminate a web proxy server from the path between a station and the switch:

  • Even if you need proxy server access enabled in order to use other applications, you can still eliminate proxy service for web access to the switch. To do so, add the IP address or DNS name of the switch to the non-proxy, or "Exceptions" list in the web browser interface you are using on the authorized station.

  • If you don't need proxy server access at all on the authorized station, then just disable the proxy server feature in the station's web browser interface.

Key Management System

The switches covered in this guide provide support for advanced routing capabilities. Security turns out to be extremely important as complex networks and the internet grow and become a part of our daily life and business. This fact forces protocol developers to improve security mechanisms employed by their protocols, which in turn becomes an extra burden for system administrators who have to set up and maintain them. One possible solution to the problem is to centralize the mechanisms used to configure and maintain security information for all routing protocols. The Key Management System (KMS) can carry this burden.

KMS is designed to configure and maintain key chains. A key chain is a set of keys with a timing mechanism for activating and deactivating individual keys. KMS provides specific instances of routing protocols with one or more Send or Accept keys that must be active at the time of a request. A *protocol instance* is usually an interface on which the protocol is running.

Feature	Default	Menu	CLI	Web
Generating a Key Chain	n/a	n/a	page 13-2	n/a
Generating a Time-Independent key	n/a	n/a	page 13-3	n/a
Generating a Time-Dependent key	n/a	n/a	page 13-4	n/a

Terminology

- **Key Chain:** A key or set of keys assigned for use by KMS-enabled protocols. A key chain may optionally contain the time to activate and deactivate a particular key.

- **Time-Independent Key:** A key that has no activate or deactivate time associated with it. This type of key does not expire, which eliminates the need for a key chain.

- **Time-Dependent key:** a key that has an activate and deactivate time associated with the Accept and Send processes. Time-Dependent keys expire, which means a key chain is needed to keep the assigned protocols supplied with keys.

- **Key Management System (KMS) Enabled Protocol:** A protocol that uses KMS to store authentication key information.

Configuring Key Chain Management

KMS-Related CLI Commands in This Section	Page
show key-chain < *chain_name* >	page 13-2
[no] key-chain *chain_name*	page 13-2
[no] key-chain *chain_name* key *Key_ID*	page 13-3

The Key Management System (KMS) has three configuration steps:

1. Create a key chain entry.

2. Assign a time-independent key or set of time-dependent keys to the Key Chain entry. The choice of key type is based on the level of security required for the protocol to which the key entry will be assigned.

3. Assign the key chain to a KMS-enabled protocol.

This procedure is protocol-dependent. For information on a specific protocol, refer to the chapter covering that protocol in the *Management and Configuration Guide* for your switch.

Creating and Deleting Key Chain Entries

To use the Key Management System (KMS), you must create one or more key chain entries. An entry can be the pointer to a single time-independent key or a chain of time-dependent keys..

Syntax: [no] key-chain < *chain_name* >

> *Generate or delete a key chain entry. Using the optional* **no** *form of the command deletes the key chain. The < chain_name > parameter can include up to 32 characters.*

show key-chain

> *Displays the current key chains on the switch and their overall status.*

For example, to generate a new key chain entry:

```
HP Switch(config)# key-chain HPN1                    Add new key chain Entry
HP-E2915-8G-PoE(config)# show key-chain              "HPN1".

 Key Chains                    Display key chain
                               entries.

  Chain Name                              Keys Active Expired
  ------------------------------- ---- ------ -------
   HPN1                                  0     0     0
```

Figure 13-1. Adding a New Key Chain Entry

After you add an entry, you can assign key(s) to it for use by a KMS-enabled protocol.

Assigning a Time-Independent Key to a Chain

A time-independent key has no Accept or Send time constraints. It is valid from boot-up until you change it. If you use a time-independent key, then it is the only key needed for a key chain entry.

Syntax: [no] key-chain < *chain_name* > key < *key_id* >

> *Generates or deletes a key in the key chain entry < chain_name >. Using the optional **no** form of the command deletes the key. The < key_id > is any number from 0-255.*

[key-string < *key_str* >]

> *This option lets you specify the key value for the protocol using the key. The < key_str > can be any string of up to 14 characters in length.*

[accept-lifetime infinite] [send-lifetime infinite]

> **accept-lifetime infinite:** *Allows packets with this key to be accepted at any time from boot-up until the key is removed.*
>
> **send-lifetime infinite:** A*llows the switch to send this key as authorization, from boot-up until the key is removed.*

show key-chain < *chain_name* >

> *Displays the detail information about the keys used in the key chain named < chain_name >.*

For example, to generate a new time-independent key for the HPN1 key chain entry:

```
HP Switch(config)# key-chain HPN1 key 1        Adds a new Time-Independent key to the
HP Switch(config)# show key-chain HPN1         "HPN1" chain.

 Chain - HPN1                                   Displays keys in the key chain
                                                entry.

  Key | Accept Start GMT  Accept Stop GMT  Send Start GMT    Send Stop GMT
  --- + ----------------- ---------------- ----------------- -------------
  1   | Bootup            Infinite         Bootup            Infinite
```

Figure 13-2. Example of Adding and Displaying a Time-Independent Key to a Key Chain Entry

Assigning Time-Dependent Keys to a Chain

A time-dependent key has Accept or Send time constraints. It is valid only during the times that are defined for the key . If a time-dependent key is used, there is usually more than one key in the key chain entry.

Syntax: [no] key-chain < *chain_name* > key < *key_id* >

> *Generates or deletes a key in the key chain entry < chain_name >. Using the optional* **no** *form of the command deletes the key. The < key_id > is any number from 0-255.*

[key-string < *key_str* >]

> *This option specifies the key value referenced by the protocol using the key. The < key_str > can be any string up to 14 characters in length.*

accept-lifetime < *mm/dd/yy* [*yy*] *hh:mm:ss* | now >

> *Specifies the* **start** *date and time of the valid period in which the switch can use this key to authenticate inbound packets.*

duration < *mm/dd/yy* [*yy*] *hh:mm:ss* | *seconds* >

> *Specifies the* **time period** *during which the switch can use this key to authenticate inbound packets. Duration is either an end date and time or the number of seconds to allow after the start date and time (which is the* **accept-lifetime** *setting).*

send-lifetime <*mm/dd/yy[yy] hh:mm:ss* | now>

> *Specifies the* **start** *date and time of the valid period in which the switch can transmit this key as authentication for outbound packets.*

duration < *mm/dd/yy[yy] hh:mm:ss* | *seconds* >

> *Specifies the **time period** during which the switch can use this key to authenticate outbound packets. Duration is either an end date and time or the number of seconds to allow after the start date and time (which is the **accept-lifetime** setting).*

show key-chain < *chain_name* >

> *Displays the detail information about the keys used in the key chain named < chain_name >.*

N o t e   Using time-dependent keys requires that all the switches have accurate, synchronized time settings. You can manually set the time or use the Time protocol feature included in the switches. For more information, refer to the chapter covering time protocols in the *Management and Configuration Guide* for your switch.

For example, to add a number of keys to the key chain entry "HPN2":

```
HP Switch(config)# key—chain HPN2 key 1 accept—lifetime 01/17/11 8:00:0   Adds a key
01/18/03 8:10:00 send—lifetime 01/17/1 8:00:00 01/18/03 8:00:00           with full
HP Switch(config)# key—chain HPN2 key 2 accept—lifetime 01/18/11 8:00:0   time and
duration 87000 send—lifetime 01/18/11 8:00:00 duration 86400             date
HP Switch(config)# key—chain HPN2 key 3 accept—lifetime 01/19/1 8:00:00   Adds a
duration 87000 send—lifetime 01/19/1 8:00:00 duration 86400              key with
HP Switch(config)# key—chain HPN2 key 4 accept—lifetime 01/20/1 8:00:0    duration
duration 87000 send—lifetime 01/20/11 8:00:00 duration 86400             expresse
HP Switch(config)# key—chain HPN2 key 5 accept—lifetime 01/21/11 8:00:0   d in
duration 87000 send—lifetime 01/21/11 8:00:00 duration 86400             seconds.
```
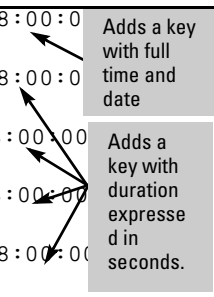
Figure 13-3. Adding Time-Dependent Keys to a Key Chain Entry

N o t e   Given transmission delays and the variations in the time value from switch to switch, it is advisable to include some flexibility in the Accept lifetime of the keys you configure. Otherwise, the switch may disregard some packets because either their key has expired while in transport or there are significant time variations between switches.

To list the result of the commands in figure 13-3:

```
HP Switch(config)# show key-chain HPN2
 Chain - HPN2

 Key | Accept Start GMT  Accept Stop GMT   Send Start GMT    Send Stop GMT
 --- + ---------------- ---------------- ---------------- ----------------
  1  | 01/17/11 08:00:00 01/18/11 08:10:00 01/17/11 08:00:00 01/18/11 08:00:00
  2  | 01/18/11 08:00:00 01/19/11 08:10:00 01/18/11 08:00:00 01/19/11 08:00:00
  3  | 01/19/11 08:00:00 01/20/11 08:10:00 01/19/11 08:00:00 01/20/11 08:00:00
  4  | 01/20/11 08:00:00 01/21/11 08:10:00 01/20/11 08:00:00 01/21/11 08:00:00
  5  | 01/21/11 08:00:00 01/22/11 08:10:00 01/21/11 08:00:00 01/22/11 08:00:00
```

Figure 13-4. Display of Time-Dependent Keys in the Key Chain Entry

You can use **show key-chain** to display the key status at the time the command is issued. Using the information from the example configuration in figures 13-3 and 13-4, if you execute **show key-chain** at 8:05 on 01/19/11, the display would appear as follows:

```
HP Switch(config)# show key-chain
 Key Chains

  Chain Name                        Keys Active Expired
  ------------------------------- ---- ------ -------
   HPN1                            1    1      0
   HPN2                            5    2      1
```

Figure 13-5. Status of Keys in Key Chain Entry "HPN2"

The "HPN1" key chain entry is a time-independent key and will not expire. "HPN2" uses time-dependent keys, which result in this data:

Expired = 1	Key 1 has expired because its lifetime ended at 8:10 on 01/18/11, the previous day.
Active = 2	Key 2 and 3 are both active for 10 minutes from 8:00 to 8:10 on 1/19/11.

Keys 4 and 5 are either not yet active or expired. The total number of keys is 5.

Index

Technology for better business outcomes

To learn more, visit www.hp.com/networking/support/networking