



Access Security Guide

2810

ProCurve Series 2810 Switches

N.10.XX

www.procurve.com



ProCurve Series 2810 Switches

July 2006

Access Security Guide

© Copyright 2001-2006 Hewlett-Packard Company, L.P.
The information contained herein is subject to change without notice.

Publication Number

5991-4734
July 2006

Applicable Products

ProCurve Switch 2810-24 (J9021A)
ProCurve Switch 2810-48 (J9022A)

Trademark Credits

Windows NT®, Windows®, and MS Windows® are US registered trademarks of Microsoft Corporation.

Software Credits

SSH on ProCurve Switches is based on the OpenSSH software toolkit. This product includes software developed by the OpenSSH Project for use in the OpenSSH Toolkit. For more information on OpenSSH, visit <http://www.openssh.com>.

SSL on ProCurve 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 <http://www.openssl.org>.

This product includes cryptographic software written by Eric Young (ey@cryptsoft.com)

This product includes software written by Tim Hudson (tjh@cryptsoft.com)

Disclaimer

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

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

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

Warranty

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

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

Contents

Product Documentation

About Your Switch Manual Set	xi
Feature Index	xii

1 Getting Started

Contents	1-1
Introduction	1-2
Overview of Access Security Features	1-2
Management Access Security Protection	1-3
General Switch Traffic Security Guidelines	1-4
Conventions	1-5
Feature Descriptions by Model	1-5
Command Syntax Statements	1-5
Command Prompts	1-6
Screen Simulations	1-6
Port Identity Examples	1-6
Sources for More Information	1-7
Need Only a Quick Start?	1-8
IP Addressing	1-8
To Set Up and Install the Switch in Your Network	1-9

2 Configuring Username and Password Security

Contents	2-1
Overview	2-2
Configuring Local Password Security	2-4
Menu: Setting Passwords	2-4
CLI: Setting Passwords and Usernames	2-5
Web: Setting Passwords and Usernames	2-6

Front-Panel Security	2-7
When Security Is Important	2-7
Front-Panel Button Functions	2-8
Configuring Front-Panel Security	2-10
Password Recovery	2-15
Password Recovery Process	2-17

3 Web and MAC Authentication

Contents	3-1
Overview	3-2
Client Options	3-3
General Features	3-4
How Web and MAC Authentication Operate	3-5
Authenticator Operation	3-5
Terminology	3-9
Operating Rules and Notes	3-10
General Setup Procedure for Web/MAC Authentication	3-12
Do These Steps Before You Configure Web/MAC Authentication ..	3-12
Additional Information for Configuring the RADIUS Server To Support MAC Authentication	3-14
Configuring the Switch To Access a RADIUS Server	3-15
Configuring Web Authentication	3-17
Overview	3-17
Configure the Switch for Web-Based Authentication	3-18
Configuring MAC Authentication on the Switch	3-22
Overview	3-22
Configure the Switch for MAC-Based Authentication	3-23
Show Status and Configuration of Web-Based Authentication	3-26
Show Status and Configuration of MAC-Based Authentication	3-27
Show Client Status	3-29

4 TACACS+ Authentication

Contents	4-1
Overview	4-2
Terminology Used in TACACS Applications:	4-3
General System Requirements	4-5
General Authentication Setup Procedure	4-5
Configuring TACACS+ on the Switch	4-8
Before You Begin	4-8
CLI Commands Described in this Section	4-9
Viewing the Switch's Current Authentication Configuration	4-9
Viewing the Switch's Current TACACS+ Server Contact Configuration	4-10
Configuring the Switch's Authentication Methods	4-11
Configuring the Switch's TACACS+ Server Access	4-15
How Authentication Operates	4-20
General Authentication Process Using a TACACS+ Server	4-20
Local Authentication Process	4-22
Using the Encryption Key	4-23
Controlling Web Browser Interface Access When Using TACACS+ Authentication	4-24
Messages Related to TACACS+ Operation	4-25
Operating Notes	4-26

5 RADIUS Authentication and Accounting

Contents	5-1
Overview	5-2
Terminology	5-3
Switch Operating Rules for RADIUS	5-4
General RADIUS Setup Procedure	5-5
Configuring the Switch for RADIUS Authentication	5-6
Outline of the Steps for Configuring RADIUS Authentication	5-7

1. Configure Authentication for the Access Methods You Want RADIUS To Protect	5-8
2. Configure the Switch To Access a RADIUS Server	5-10
3. Configure the Switch's Global RADIUS Parameters	5-12
Local Authentication Process	5-16
Controlling Web Browser Interface Access When Using RADIUS Authentication	5-17
Configuring RADIUS Accounting	5-17
Operating Rules for RADIUS Accounting	5-19
Steps for Configuring RADIUS Accounting	5-19
Viewing RADIUS Statistics	5-25
General RADIUS Statistics	5-25
RADIUS Authentication Statistics	5-27
RADIUS Accounting Statistics	5-28
Changing RADIUS-Server Access Order	5-30
Messages Related to RADIUS Operation	5-32

6 Configuring Secure Shell (SSH)

Contents	6-1
Overview	6-2
Terminology	6-4
Prerequisite for Using SSH	6-5
Public Key Formats	6-5
Steps for Configuring and Using SSH for Switch and Client Authentication	6-6
General Operating Rules and Notes	6-8
Configuring the Switch for SSH Operation	6-9
1. Assign Local Login (Operator) and Enable (Manager) Password ..	6-9
2. Generate the Switch's Public and Private Key Pair	6-10
3. Provide the Switch's Public Key to Clients	6-13
4. Enable SSH on the Switch and Anticipate SSH Client Contact Behavior	6-15
5. Configure the Switch for SSH Authentication	6-18

6. Use an SSH Client To Access the Switch	6-22
Further Information on SSH Client Public-Key Authentication	6-22
Messages Related to SSH Operation	6-28

7 Configuring Secure Socket Layer (SSL)

Contents	7-1
Overview	7-2
Terminology	7-3
Prerequisite for Using SSL	7-5
Steps for Configuring and Using SSL for Switch and Client Authentication	7-5
General Operating Rules and Notes	7-6
1. Assign Local Login (Operator) and Enable (Manager) Password	7-7
2. Generate the Switch's Server Host Certificate	7-9
3. Enable SSL on the Switch and Anticipate SSL Browser Contact Behavior	7-17
Common Errors in SSL Setup	7-21

8 Configuring Port-Based Access Control (802.1X)

Contents	8-1
Overview	8-3
Why Use Port-Based Access Control?	8-3
General Features	8-3
How 802.1X Operates	8-6
Authenticator Operation	8-6
Switch-Port Supplicant Operation	8-7
Terminology	8-8
General Operating Rules and Notes	8-10
General Setup Procedure for Port-Based Access Control (802.1X)	8-12
Do These Steps Before You Configure 802.1X Operation	8-12
Overview: Configuring 802.1X Authentication on the Switch	8-13
Configuring Switch Ports as 802.1X Authenticators	8-15

1. Enable 802.1X Authentication on Selected Ports	8-15
3. Configure the 802.1X Authentication Method	8-19
4. Enter the RADIUS Host IP Address(es)	8-20
5. Enable 802.1X Authentication on the Switch	8-20
802.1X Open VLAN Mode	8-21
Introduction	8-21
Use Models for 802.1X Open VLAN Modes	8-22
Operating Rules for Authorized-Client and Unauthorized-Client VLANs	8-25
Setting Up and Configuring 802.1X Open VLAN Mode	8-27
802.1X Open VLAN Operating Notes	8-31
Option For Authenticator Ports: Configure Port-Security To Allow Only 802.1X Devices	8-32
Configuring Switch Ports To Operate As Supplicants for 802.1X Connections to Other Switches	8-34
Displaying 802.1X Configuration, Statistics, and Counters	8-38
Show Commands for Port-Access Authenticator	8-38
Viewing 802.1X Open VLAN Mode Status	8-40
Show Commands for Port-Access Supplicant	8-43
How RADIUS/802.1X Authentication Affects VLAN Operation	8-44
Messages Related to 802.1X Operation	8-48

9 Configuring and Monitoring Port Security

Contents	9-1
Overview	9-2
Basic Operation	9-2
Blocking Unauthorized Traffic	9-3
Trunk Group Exclusion	9-4
Planning Port Security	9-5
Port Security Command Options and Operation	9-6
Retention of Static MAC Addresses	9-10
Displaying Current Port Security Settings	9-10
Configuring Port Security	9-12
MAC Lockdown	9-17

Differences Between MAC Lockdown and Port Security	9-19
Deploying MAC Lockdown	9-21
MAC Lockout	9-25
Port Security and MAC Lockout	9-27
Web: Displaying and Configuring Port Security Features	9-27
Reading Intrusion Alerts and Resetting Alert Flags	9-28
Notice of Security Violations	9-28
How the Intrusion Log Operates	9-29
Keeping the Intrusion Log Current by Resetting Alert Flags	9-29
Using the Event Log To Find Intrusion Alerts	9-35
Web: Checking for Intrusions, Listing Intrusion Alerts, and Resetting Alert Flags	9-36
Operating Notes for Port Security	9-36

10 Traffic/Security Filters

Contents	10-1
Overview	10-2
Using Source-Port Filters	10-3
Operating Rules for Source-Port Filters	10-3
Configuring a Source-Port Filter	10-4
Viewing a Source-Port Filter	10-7
Filter Indexing	10-8
Editing a Source-Port Filter	10-9
Using Named Source-Port Filters	10-9

11 Using Authorized IP Managers

Contents	11-1
Overview	11-2
Configuration Options	11-3
Access Levels	11-3
Defining Authorized Management Stations	11-4
Overview of IP Mask Operation	11-4
Menu: Viewing and Configuring IP Authorized Managers	11-5

CLI: Viewing and Configuring Authorized IP Managers	11-6
Web: Configuring IP Authorized Managers	11-9
Building IP Masks	11-9
Configuring One Station Per Authorized Manager IP Entry	11-9
Configuring Multiple Stations Per Authorized Manager IP Entry ..	11-10
Additional Examples for Authorizing Multiple Stations	11-12
Operating Notes	11-12

Index

Product Documentation

About Your Switch Manual Set

The switch manual set includes the following:

- **Read Me First** - a printed guide shipped with your switch. Provides software update information, product notes, and other information.
- **Installation and Getting Started Guide** - a printed guide shipped with your switch. This guide explains how to prepare for and perform the physical installation and connection to your network.
- **Management and Configuration Guide** - a PDF file on the ProCurve Networking website. This guide describes how to configure, manage, and monitor basic switch operation.
- **Advanced Traffic Management Guide** - a PDF file on the ProCurve Networking website. This guide explains the configuration and operation of traffic management features such as spanning tree and VLANs.
- **Access Security Guide** - a PDF file on the ProCurve Networking website. This guide explains the configuration and operation of access security and user authentication features on the switch.
- **Release Notes** - posted on the ProCurve web site to provide information on software updates. The release notes describe new features, fixes, and enhancements that become available between revisions of the above guides.

Note

For the latest version of all ProCurve switch documentation, including release notes covering recently added features, visit the ProCurve Networking website at <http://www.procurve.com>. Click on **Technical support**, and then click on **Product manuals**.

Feature Index

For the manual set supporting your switch model, the following feature index indicates which manual to consult for information on a given software feature.

Feature	Management and Configuration	Advanced Traffic Management	Access Security Guide
802.1Q VLAN Tagging	-	X	-
802.1X Port-Based Priority	X	-	-
Authentication	-	-	X
Authorized IP Managers	-	-	X
Config File	X	-	-
Copy Command	X	-	-
Debug	X	-	-
DHCP Configuration	-	X	-
DHCP/Bootp Operation	X	-	-
Diagnostic Tools	X	-	-
Downloading Software	X	-	-
Event Log	X	-	-
Factory Default Settings	X	-	-
File Management	X	-	-
File Transfers	X	-	-
GVRP	-	X	-
IGMP	-	X	-
Interface Access (Telnet, Console/Serial, Web)	X	-	-
IP Addressing	X	-	-
LACP	X	-	-
Link	X	-	-

Feature	Management and Configuration	Advanced Traffic Management	Access Security Guide
LLDP	X	-	-
MAC Address Management	X	-	-
MAC Lockdown	-	-	X
MAC Lockout	-	-	X
MAC-based Authentication	-	-	X
Monitoring and Analysis	X	-	-
Multicast Filtering	-	X	-
Network Management Applications (LLDP, SNMP)	X	-	-
Passwords	-	-	X
Ping	X	-	-
Port Configuration	X	-	-
Port Security	-	-	X
Port Status	X	-	-
Port Trunking (LACP)	X	-	-
Port-Based Access Control	-	-	X
Port-Based Priority (802.1Q)	X	-	-
Quality of Service (QoS)	-	X	-
RADIUS Authentication and Accounting	-	-	X
Secure Copy	X	-	-
SFTP	X	-	-
SNMP	X	-	-
Software Downloads (SCP/SFTP, TFTP, Xmodem)	X	-	-
Source-Port Filters	-	-	X
Spanning Tree (MSTP)	-	X	-
SSH (Secure Shell) Encryption	-	-	X
SSL (Secure Socket Layer)	-	-	X

Product Documentation

Feature	Management and Configuration	Advanced Traffic Management	Access Security Guide
Stack Management (Stacking)	-	X	-
Syslog	X	-	-
System Information	X	-	-
TACACS+ Authentication	-	-	X
Telnet Access	X	-	-
TFTP	X	-	-
Time Protocols (TimeP, SNTP)	X	-	-
Traffic/Security Filters	-	-	X
Troubleshooting	X	-	-
VLANs	-	X	-
Web-based Authentication	-	-	X
Xmodem	X	-	-

Getting Started

Contents

Introduction	1-2
Overview of Access Security Features	1-2
Management Access Security Protection	1-3
General Switch Traffic Security Guidelines	1-4
Conventions	1-5
Feature Descriptions by Model	1-5
Command Syntax Statements	1-5
Command Prompts	1-5
Screen Simulations	1-6
Port Identity Examples	1-6
Sources for More Information	1-6
Need Only a Quick Start?	1-8
IP Addressing	1-8
To Set Up and Install the Switch in Your Network	1-8

Introduction

This *Access Security Guide* describes how to use ProCurve's switch security features to protect access to your switch. This guide is intended to support the following switches:

- ProCurve Switch 2810

For an overview of other product documentation for the above switches, refer to "Product Documentation" on page xi.

You can download a copy from the ProCurve Networking website, <http://www.procurve.com>.

Overview of Access Security Features

The access security features covered in this guide include:

- **Local Manager and Operator Passwords** (page 2-1): Control access and privileges for the CLI, menu, and web browser interfaces.
- **TACACS+ Authentication** (page 4-1): Uses an authentication application on a server to allow or deny access to a switch.
- **RADIUS Authentication and Accounting** (page 5-1): Like TACACS+, uses an authentication application on a central server to allow or deny access to the switch. RADIUS also provides accounting services for sending data about user activity and system events to a RADIUS server.
- **Secure Shell (SSH) Authentication** (page 6-1): Provides encrypted paths for remote access to switch management functions.
- **Secure Socket Layer (SSL)** (page 7-1): Provides remote web access to the switch via encrypted authentication paths between the switch and management station clients capable of SSL/TLS operation.

- **Port-Based Access Control (802.1X)** (page 8-1): On point-to-point connections, enables the switch to allow or deny traffic between a port and an 802.1X-aware device (supplicant) attempting to access the switch. Also enables the switch to operate as a supplicant for connections to other 802.1X-aware switches.
- **Port Security** (page 9-1): Enables a switch port to maintain a unique list of MAC addresses defining which specific devices are allowed to access the network through that port. Also enables a port to detect, prevent, and log access attempts by unauthorized devices.
- **Traffic/Security Filters** (page 10-1): Source-Port filtering enhances in-band security by enabling outbound destination ports on the switch to forward or drop traffic from designated source ports (within the same VLAN).
- **Authorized IP Managers** (page 11-1): Allows access to the switch by a networked device having an IP address previously configured in the switch as “authorized”.

Management Access Security Protection

In considering management access security for your switch, there are two key areas to protect:

- Unauthorized client access to switch management features
- Unauthorized client access to the network.

Table 1-1 on page 1-4 provides an overview of the type of protection offered by each switch security feature.

Note

ProCurve recommends that you use local passwords together with your switch’s other security features to provide a more comprehensive security fabric than if you use only local passwords.

Table 1-1. Management Access Security Protection

Security Feature	Offers Protection Against Unauthorized Client Access to Switch Management Features					Offers Protection Against Unauthorized Client Access to the Network
	Connection	Telnet	SNMP (Net Mgmt)	Web Browser	SSH Client	
Local Manager and Operator Usernames and Passwords ¹	PtP:	Yes	<i>No</i>	Yes	Yes	<i>No</i>
	Remote:	Yes	<i>No</i>	Yes	Yes	<i>No</i>
TACACS+	PtP:	Yes	<i>No</i>	<i>No</i>	Yes	<i>No</i>
	Remote:	Yes	<i>No</i>	<i>No</i>	Yes	<i>No</i>
RADIUS	PtP:	Yes	<i>No</i>	<i>No</i>	Yes	<i>No</i>
	Remote:	Yes	<i>No</i>	<i>No</i>	Yes	<i>No</i>
SSH	PtP:	Yes	<i>No</i>	<i>No</i>	Yes	<i>No</i>
	Remote:	Yes	<i>No</i>	<i>No</i>	Yes	<i>No</i>
SSL	PtP:	<i>No</i>	<i>No</i>	Yes	<i>No</i>	<i>No</i>
	Remote:	<i>No</i>	<i>No</i>	Yes	<i>No</i>	<i>No</i>
Port-Based Access Control (802.1X)	PtP:	Yes	Yes	Yes	Yes	Yes
	Remote:	<i>No</i>	<i>No</i>	<i>No</i>	<i>No</i>	<i>No</i>
Port Security (MAC address)	PtP:	Yes	Yes	Yes	Yes	Yes
	Remote:	Yes	Yes	Yes	Yes	Yes
Authorized IP Managers	PtP:	Yes	Yes	Yes	Yes	<i>No</i>
	Remote:	Yes	Yes	Yes	Yes	<i>No</i>

General Switch Traffic Security Guidelines

Where the switch is running multiple security options, it implements network traffic security based on the OSI (Open Systems Interconnection model) precedence of the individual options, from the lowest to the highest. The following list shows the order in which the switch implements configured security features on traffic moving through a given port.

1. Disabled/Enabled physical port
2. MAC lockout (applies to all ports on the switch)
3. MAC lockdown
4. Port security
5. Authorized IP Managers
6. Application features at higher levels in the OSI model, such as SSH

(The above list does not address the mutually exclusive relationship that exists among some security features.)

Conventions

This guide uses the following conventions for command syntax and displayed information.

Feature Descriptions by Model

In cases where a software feature is not available in all of the switch models covered by this guide, the section heading specifically indicates which product or product series offer the feature.

For example (the switch model is highlighted here in ***bold italics***):

“Web and MAC Authentication for the ***2810 Switch***”.

Command Syntax Statements

Syntax: `aaa port-access authenticator < port-list >`
 [control < authorized | auto | unauthorized >]

- Vertical bars (|) separate alternative, mutually exclusive elements.
- Square brackets ([]) indicate optional elements.
- Braces (< >) enclose required elements.
- Braces within square brackets ([< >]) indicate a required element within an optional choice.
- Boldface indicates use of a CLI command, part of a CLI command syntax, or other displayed element in general text. For example:

“Use the **copy tftp** command to download the key from a TFTP server.”

- Italics indicate variables for which you must supply a value when executing the command. For example, in this command syntax, **< port-list >** indicates that you must provide one or more port numbers:

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

Command Prompts

In the default configuration, your switch displays the following CLI prompt:

```
ProCurve Switch 2810#
```

To simplify recognition, this guide uses `ProCurve` to represent command prompts for all models. For example:

```
ProCurve#
```

(You can use the **hostname** command to change the text in the CLI prompt.)

Screen Simulations

Figures containing simulated screen text and command output look like this:

```
ProCurve(config)# show version
Image stamp:      /sw/code/build/bass(ppne_swt)
                  Mar 17 2006 11:44:02
                  N.10.XX
                  2624
Boot Image:      Primary
Build Options:   QA
Watchdog:        ENABLED
```

Figure 1-1. Example of a Figure Showing a Simulated Screen

In some cases, brief command-output sequences appear outside of a numbered figure. For example:

```
ProCurve(config)# ip default-gateway 18.28.152.1/24
ProCurve(config)# vlan 1 ip address 18.28.36.152/24
ProCurve(config)# vlan 1 ip igmp
```

Port Identity Examples

This guide describes software applicable to both chassis-based and stackable ProCurve switches. Where port identities are needed in an example, this guide uses the chassis-based port identity system, such as “A1”, “B3 - B5”, “C7”, etc. However, unless otherwise noted, such examples apply equally to the stackable switches, which for port identities typically use only numbers, such as “1”, “3-5”, “15”, etc.

Sources for More Information

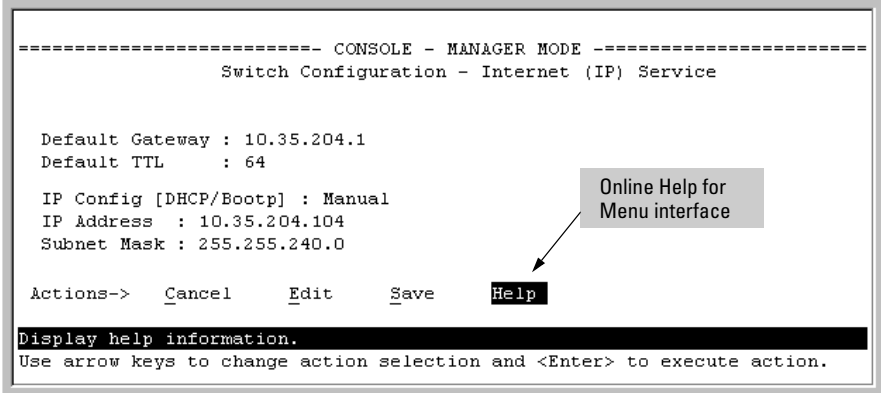
For additional information about switch operation and features not covered in this guide, consult the following sources:

- For information on which product manual to consult on a given software feature, refer to “Product Documentation” on page xi.

Note

For the latest version of all ProCurve switch documentation, including release notes covering recently added features, visit the ProCurve Networking website at <http://www.procurve.com>. Click on **Technical support**, and then click on **Product manuals**.

- For information on specific parameters in the menu interface, refer to the online help provided in the interface. For example:



```
----- CONSOLE - MANAGER MODE -----  
Switch Configuration - Internet (IP) Service  
  
Default Gateway : 10.35.204.1  
Default TTL      : 64  
  
IP Config [DHCP/Bootp] : Manual  
IP Address   : 10.35.204.104  
Subnet Mask  : 255.255.240.0  
  
Actions->  _Cancel      _Edit      _Save      Help  
  
Display help information.  
Use arrow keys to change action selection and <Enter> to execute action.
```

Figure 1-2. Getting Help in the Menu Interface

- For information on a specific command in the CLI, type the command name followed by “help”. For example:

```
ProCurve# write help
Usage: write <memory|terminal>

Description: View or save the running configuration of the switch.

write terminal - displays the running configuration of the
                switch on the terminal
write memory   - saves the running configuration of the
                switch to flash. The saved configuration
                becomes the boot-up configuration of the switch
                the next time it is booted.
```

Figure 1-3. Getting Help in the CLI

- For information on specific features in the Web browser interface, use the online help. For more information, refer to the *Management and Configuration Guide* for your switch.
- For further information on ProCurve Networking switch technology, visit the ProCurve website at:

<http://www.procurve.com>

Need Only a Quick Start?

IP Addressing

If you just want to give the switch an IP address so that it can communicate on your network, or if you are not using multiple VLANs, ProCurve recommends that you use the Switch Setup screen to quickly configure IP addressing. To do so, do one of the following:

- Enter **setup** at the CLI Manager level prompt.
ProCurve# setup
- In the Main Menu of the Menu interface, select
8. Run Setup

For more on using the Switch Setup screen, see the *Quick Installation Guide* you received with the switch.

To Set Up and Install the Switch in Your Network

Important!

Use the *Quick Installation Guide* shipped with your switch for the following:

- Notes, cautions, and warnings related to installing and using the switch
- Instructions for physically installing the switch in your network
- Quickly assigning an IP address and subnet mask, setting a Manager password, and (optionally) configuring other basic features.
- Interpreting LED behavior.

For the latest version of the *Installation and Getting Started Guide* and other documentation for your switch, visit the ProCurve Networking Web site. (Refer to “Product Documentation” on page xi of this guide for further details.)

— *This page is intentionally unused.* —

Configuring Username and Password Security

Contents

Overview	2-2
Configuring Local Password Security	2-4
Menu: Setting Passwords	2-4
CLI: Setting Passwords and Usernames	2-5
Web: Setting Passwords and Usernames	2-6
Front-Panel Security	2-7
When Security Is Important	2-7
Front-Panel Button Functions	2-8
Configuring Front-Panel Security	2-10
Password Recovery	2-15
Password Recovery Process	2-17

Overview

Feature	Default	Menu	CLI	Web
Set Usernames	none	—	—	page 2-6
Set a Password	none	page 2-4	page 2-5	page 2-6
Delete Password Protection	n/a	page 2-4	page 2-6	page 2-6
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.

Level	Actions Permitted
Manager:	Access to all console interface areas. <i>This is the default level.</i> That is, if a Manager password has <i>not</i> been set prior to starting the current console session, then anyone having access to the console can access any area of the console interface.
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 *Management and Configuration 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 username requires either the CLI or the web browser interface.

1. From the Main Menu select:

3. Console Passwords




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 password 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 > [user-name ASCII-STR]
[no] password < all >

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

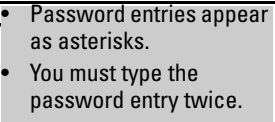


Figure 2-2. Example of Configuring Manager and Operator Passwords

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:

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

Web: Setting Passwords and Usernames

In the web browser interface you can enter passwords and (optional) usernames.

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.

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 2-9.)
- 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 2-9) 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 2-9) 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 2-17.) (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 2-9) 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.

```
ProCurve(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 2-9.

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:

```
ProCurve(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
ProCurve(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 enables 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 2-9. 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).

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

ProCurve(config)# no front-panel-security password-clear reset-on-clear
ProCurve(config)# show front-panel-security
Clear Password      - Enabled
Reset-on-clear      - Disabled
Factory Reset       - Enabled
Password Recovery   - Enabled
```

Shows password-clear disabled.

Enables password-clear, with reset-on-clear disabled by the "no" statement at the beginning of the command.

Shows password-clear enabled, with reset-on-clear disabled.

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 2-9 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 2-9:

- *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.*


```

ProCurve(config)# no front-panel-security factory-reset
                    **** CAUTION ****
Disabling the factory reset option prevents switch configuration 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
ProCurve(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 your ProCurve 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.

Caution

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 2-9) 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 the 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.

```
ProCurve(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
ProCurve(config)# _
```

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 ProCurve.

Note

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 2-9. 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 ProCurve Customer Care Center for further assistance. Using the switch's MAC address, the ProCurve 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.

Note

The alternate password provided by the ProCurve 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 ProCurve Customer Care Center.

Web and MAC Authentication

Contents

Overview	3-2
Client Options	3-3
General Features	3-4
How Web and MAC Authentication Operate	3-5
Authenticator Operation	3-5
Terminology	3-9
Operating Rules and Notes	3-10
General Setup Procedure for Web/MAC Authentication	3-12
Do These Steps Before You Configure Web/MAC Authentication ..	3-12
Additional Information for Configuring the RADIUS Server To Support MAC Authentication	3-13
Configuring the Switch To Access a RADIUS Server	3-14
Configuring Web Authentication	3-17
Overview	3-17
Configure the Switch for Web-Based Authentication	3-18
Configuring MAC Authentication on the Switch	3-22
Overview	3-22
Configure the Switch for MAC-Based Authentication	3-23
Show Status and Configuration of Web-Based Authentication	3-26
Show Status and Configuration of MAC-Based Authentication	3-27
Show Client Status	3-29

Overview

Feature	Default	Menu	CLI	Web
Configure Web Authentication	n/a	—	3-17	—
Configure MAC Authentication	n/a	—	3-22	—
Display Web Authentication Status and Configuration	n/a	—	3-26	—
Display MAC Authentication Status and Configuration	n/a	—	3-27	—

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.

Note

You can configure only one authentication type on a port. This means that Web authentication, MAC authentication, 802.1X, MAC lockdown, MAC lock-out, and port-security are mutually exclusive on a given port. Also, LACP must be disabled on ports configured for any of these authentication methods.

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. However, where all clients can operate in the same VLAN, the switch allows up to 32 simultaneous clients per port. (In applications where you want the switch to simultaneously support multiple client sessions in different VLANs, design your system so that such clients will use different switch ports.)

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 the CHAP 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.

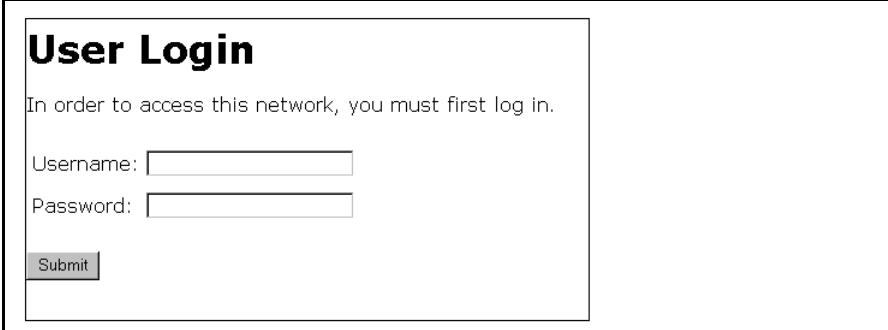
How Web and MAC Authentication Operate

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:

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:

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-based 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.

Web and MAC Authentication

How Web and MAC Authentication Operate

-
-
-
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.

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 ProCurve 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. ProCurve 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 3-21.

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.

Operating Rules and Notes

- You can configure one type of authentication on a port. That is, the following authentication types are *mutually exclusive* on a given port:
 - Web Authentication
 - MAC Authentication
 - 802.1X
- 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.

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, ProCurve 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 3-12.)
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 3-15.)

Additional Information for 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 four format options:

aabbccddeeff (the default format)

aabbcc-ddeeff

aa-bb-cc-dd-ee-ff

aa:bb:cc:dd:ee:ff

Note on MAC Addresses

Letters in MAC addresses must be in lowercase.

- 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 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 configure up to three RADIUS server addresses. The switch uses the first server it successfully accesses. (Refer to “RADIUS Authentication and Accounting” 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

Web and MAC Authentication

Configuring the Switch To Access a RADIUS Server

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, 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'

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

Status and Counters - General RADIUS Information

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

  Server IP Addr   Auth  Acct
  -----
  192.168.32.11   1812  1813  2Pzo22

ProCurve(config)#
```

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. ProCurve 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-Based Authentication

Command	Page
Configuration Level	
aaa port-access web-based dhcp-addr	3-18
aaa port-access web-based dhcp-lease	3-18
[no] aaa port-access web-based [e] <port-list>	3-19
[auth-vid]	3-19
[client-limit]	3-19
[client-moves]	3-19
[logoff-period]	3-20
[max-requests]	3-20
[max-retries]	3-20
[quiet-period]	3-20
[reauth-period]	3-20
[reauthenticate]	3-20
[redirect-url]	3-21
[server-timeout]	3-21
[ssl-login]	3-21
[unauth-vid]	3-22

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>]
no 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 <1-32>]

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`. ProCurve 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	3-23
[no] aaa port-access mac-based [e] < port-list >	3-23
[addr-limit]	3-24
[addr-moves]	3-24
[auth-vid]	3-24
[logoff-period]	3-24
[max-requests]	3-24
[quiet-period]	3-25
[reauth-period]	3-25
[reauthenticate]	3-25
[server-timeout]	3-25
[unauth-vid]	3-25

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.

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 <1-32>]
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. Depending 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 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)*

Show Status and Configuration of Web-Based Authentication

Command	Page
show port-access [<i>port-list</i>] web-based	3-26
[clients]	3-26
[config]	3-26
[config [auth-server]]	3-27
[config [web-server]]	3-27
show port-access <i>port-list</i> web-based config detail	3-27

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 Authentication 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 [<i>port-list</i>] mac-based	3-27
[clients]	3-28
[config]	3-28
[config [auth-server]]	3-28
show port-access <i>port-list</i> mac-based config detail	3-28

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.

Web and MAC Authentication

Show Status and Configuration of MAC-Based Authentication

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	<ol style="list-style-type: none"> 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	<ol style="list-style-type: none"> 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.

— This page is intentionally unused. —

TACACS+ Authentication

Contents

Overview	4-2
Terminology Used in TACACS Applications:	4-3
General System Requirements	4-5
General Authentication Setup Procedure	4-5
Configuring TACACS+ on the Switch	4-8
Before You Begin	4-8
CLI Commands Described in this Section	4-9
Viewing the Switch's Current Authentication Configuration	4-9
Viewing the Switch's Current TACACS+ Server Contact Configuration	4-10
Configuring the Switch's Authentication Methods	4-11
Configuring the Switch's TACACS+ Server Access	4-15
How Authentication Operates	4-20
General Authentication Process Using a TACACS+ Server	4-20
Local Authentication Process	4-22
Using the Encryption Key	4-23
Controlling Web Browser Interface Access When Using TACACS+ Authentication	4-24
Messages Related to TACACS+ Operation	4-25
Operating Notes	4-25

Overview

Feature	Default	Menu	CLI	Web
view the switch's authentication configuration	n/a	—	page 4-9	—
view the switch's TACACS+ server contact configuration	n/a	—	page 4-10	—
configure the switch's authentication methods	disabled	—	page 4-11	—
configure the switch to contact TACACS+ server(s)	disabled	—	page 4-15	—

TACACS+ authentication enables you to use a central server to allow or deny access to the switch (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).

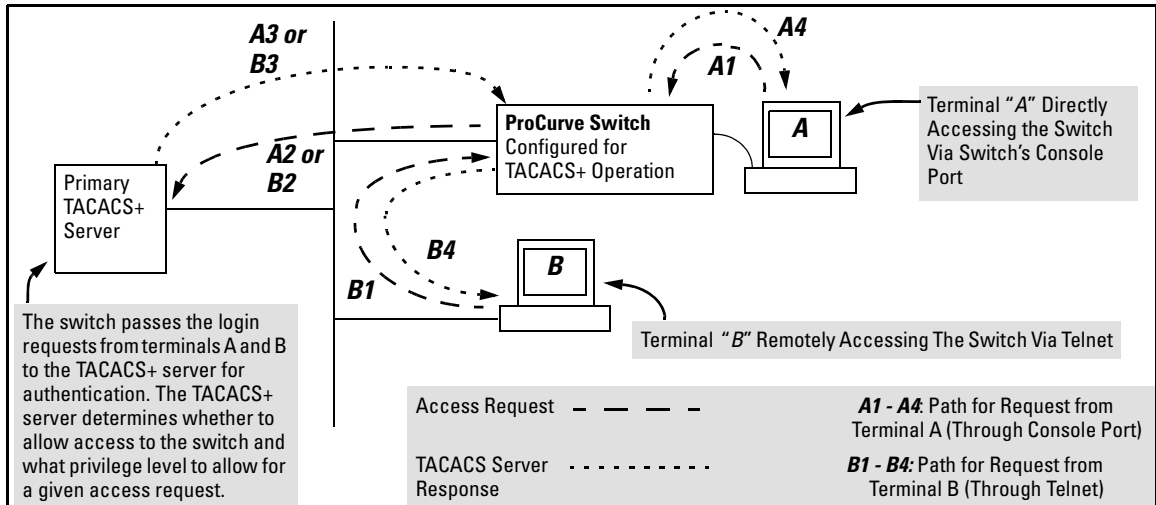


Figure 4-1. Example of TACACS+ Operation

TACACS+ in the switch 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 authentica-

tion 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.

Notes

The software does not support TACACS+ authorization or accounting services.

TACACS+ does not affect web browser interface access. See “Controlling Web Browser Interface Access” on page 4-24.

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 when TACACS+ is enabled on the switch (that is, when the switch is TACACS-aware).
- **TACACS+ Server:** The server or management station configured as an access control server for TACACS-enabled devices. To use TACACS+ with the switch 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 “Configuring Username and Password Security” on page 2-1.)
- **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, ProCurve recommends that you thoroughly test all TACACS+ configurations used in your network.

TACACS-aware ProCurve switches include the capability of configuring multiple backup TACACS+ servers. ProCurve 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 4-24.

General Authentication Setup Procedure

It is important to test the TACACS+ service before fully implementing it. Depending on the process and parameter settings you use to set up and test TACACS+ authentication in your network, you could accidentally lock all users, including yourself, out of access to a switch. While recovery is simple, it may pose an inconvenience that can be avoided. To prevent an unintentional lockout on a 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 Troubleshooting 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 a switch. (Refer to the documentation provided with the TACACS+ server software.) This includes knowing whether you need to configure an encryption key. (See “Using the Encryption Key” on page 4-23.)
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, ProCurve 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 configuration, check Telnet access again. If Telnet access still fails, check the

configuration in your TACACS+ server application for mis-configurations or missing data that could affect the server's interoperability 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 memory.

Configuring TACACS+ on the Switch

Before You Begin

If you are new to TACACS+ authentication, ProCurve recommends that you read the "General Authentication Setup Procedure" on page 4-5 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	4-9
show tacacs	4-10
aaa authentication	pages 4-11 through 4-14
console	
Telnet	
num-attempts <1-10 >	
tacacs-server	pages 4-15
host < ip-addr >	pages 4-15
key	4-19
timeout < 1-255 >	4-20

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.

```

ProCurve> 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:

```
ProCurve # show tacacs
Status and Counters - TACACS Information
Timeout : 5
Encryption Key : paris-1
Server IP Addr  Opens  Closes  Aborts  Errors  Pkts Rx  Pkts Tx
-----
10.30.248.100   0       0       0       0       0       0
10.30.248.156   0       0       0       0       0       0
10.30.248.105   0       0       0       0       0       0
```

The diagram shows three labels on the left side of the table: 'First-Choice TACACS+ Server' with an arrow pointing to the first IP address (10.30.248.100), 'Second-Choice TACACS+ Server' with an arrow pointing to the second IP address (10.30.248.156), and 'Third-Choice TACACS+ Server' with an arrow pointing to the third IP address (10.30.248.105).

Figure 4-3. Example of the Switch's TACACS+ Configuration Listing

Configuring the Switch's Authentication Methods

The **aaa authentication** command configures the access control for console port and Telnet access to the switch. That is, for both access methods, **aaa authentication** specifies whether to use a TACACS+ server or the switch's local authentication, or (for some secondary scenarios) no authentication (meaning that if the primary method fails, authentication is denied). This command also reconfigures the number of access attempts to allow in a session if the first attempt uses an incorrect username/password pair.

Syntax: aaa authentication

< console | telnet >

Selects either console (serial port) or Telnet access for configuration.

< enable | login >

Selects either the Manager (enable) or Operator (login) access level.

< 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. (Refer to "RADIUS Authentication and Accounting" on page 5-1.)

[< local | none >]

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

Note: *If the primary authentication is local, you can't use local for the secondary authentication access method.*

aaa authentication num-attempts < 1-10 >

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

Table 4-1. AAA Authentication Parameters

Name	Default	Range	Function
console - or - telnet	n/a	n/a	Specifies whether the command is configuring authentication for the console port or Telnet access method for the switch.
enable - or - login	n/a	n/a	Specifies the privilege level for the access method being configured. login: Operator (read-only) privileges enable: Manager (read-write) privileges
local - or - tacacs	local	n/a	Specifies the primary method of authentication for the access method being configured. local: Use the username/password pair configured locally in the switch for the privilege level being configured tacacs: Use a TACACS+ server.
local - or - none	none	n/a	Specifies the secondary (backup) type of authentication being configured. local: The username/password pair configured locally in the switch for the privilege level being configured. Cannot be used if the primary authentication is local . none: No secondary type of authentication for the specified method/privilege path. (Available only if the primary method of authentication for the access being configured is local.) Note: If you do not specify this parameter in the command line, the switch automatically assigns the secondary method as follows: <ul style="list-style-type: none">• If the primary method is tacacs, the only secondary method is local.• If the primary method is local, the only secondary method is none.
num-attempts	3	1 - 10	In a given session, specifies how many tries at entering the correct username/ password pair are allowed before access is denied and the session terminated.

As shown in the 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.

*When “local” is the primary option, you can also select “local” as the secondary option. However, in this case, a secondary “local” is meaningless because the switch has only one local level of username/password protection.

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.

TACACS+ Authentication

Configuring TACACS+ on the Switch

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.**

```
ProCurve (config)# aaa authentication console login tacacs local
```

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

```
ProCurve (config)# aaa authentication console enable tacacs local
```

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

```
ProCurve (config)# aaa authentication Telnet login tacacs local
```

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

```
ProCurve (config)# aaa authentication telnet enable tacacs local
```

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

```
ProCurve (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 Address 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 4-11.)

Note

As described under “General Authentication Setup Procedure” on page 4-5, ProCurve 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 on
Encryption Keys**

Encryption keys configured in the switch must exactly match the 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”.

Table 4-3. Details on Configuring TACACS Servers and Keys

Name	Default	Range
tacacs-server host <ip-addr>	none	n/a

This command specifies the IP address of a device running a TACACS+ server application. Optionally, it 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 4-23 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 4-10.)

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

1. When there are no TACACS+ servers configured, entering a server IP address makes that server the first-choice TACACS+ server.
 2. When there is one TACACS+ server already configured, entering another server IP address makes that server the second-choice (backup) TACACS+ server.
 3. When there are two TACACS+ servers already configured, entering another server IP address makes that server the third-choice (backup) TACACS+ server.
- The above position assignments are fixed. Thus, if you remove one server and replace it with another, the new server assumes the priority position that the removed server had. For example, suppose you configured three servers, A, B, and C, configured in order:
First-Choice:A
Second-Choice:B
Third-Choice: C
 - If you removed server B and then entered server X, the TACACS+ server order of priority would be:
First-Choice:A
Second-Choice:X
Third-Choice: C
 - If there are two or more vacant slots in the TACACS+ server priority list and you enter a new IP address, the new address will take the vacant slot with the highest priority. Thus, if A, B, and C are configured as above and you (1) remove A and B, and (2) enter X and Y (in that order), then the new TACACS+ server priority list would be X, Y, and C.
 - The easiest way to change the order of the TACACS+ servers in the priority list is to remove all server addresses in the list and then re-enter them in order, with the new first-choice server address first, and so on.

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

See also “General Authentication Process Using a TACACS+ Server” on page 4-20.

Name	Default	Range
[key <key-string>]	none (null)	n/a
<p>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.)</p> <p>For more on the encryption key, see “Using the Encryption Key” on page 4-23 and the documentation provided with your TACACS+ server application.</p>		
timeout <1 - 255>	5 sec	1 - 255 sec
<p>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).</p>		

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:

```

ProCurve# show tacacs
  Status and Counters - TACACS Information
  Timeout : 5
  Encryption Key :
  Server IP Addr  Opens  Closes  Aborts  Errors  Pkts Rx  Pkts Tx
  -----
  (10.28.227.15)  0       0       0       0       0       0
  10.28.227.10   0       0       0       0       0       0
    
```

Figure 4-4. 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:

The "10" server is now the "first-choice" TACACS+ authentication device.

```
ProCurve# show tacacs
Status and Counters - TACACS Information
Timeout : 5
Encryption Key :
```

Server IP	Addr	Opens	Closes	Aborts	Errors	Pkts Rx	Pkts Tx
10.28.227.10		0	0	0	0	0	0
10.28.227.15		0	0	0	0	0	0

Figure 4-5. 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:

```
ProCurve(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 4-23.)

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

```
ProCurve(config) tacacs-server key north01
```

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

```
ProCurve(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:

```
ProCurve(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:

```
ProCurve(config)# tacacs-server host 10.28.227.104
```

Note

The `show tacacs` command lists the global encryption key, if configured. However, to view any configured per-server encryption keys, you must use **show config** 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:

```
ProCurve(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 4-6. Using a TACACS+ Server for Authentication

Using figure 4-6, 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 4-22.)
 - 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 4-13.)

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:

```
ProCurve(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:

```
ProCurve(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.
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.)

RADIUS Authentication and Accounting

Contents

Overview	5-2
Terminology	5-3
Switch Operating Rules for RADIUS	5-4
General RADIUS Setup Procedure	5-5
Configuring the Switch for RADIUS Authentication	5-6
Outline of the Steps for Configuring RADIUS Authentication	5-7
1. Configure Authentication for the Access Methods You Want RADIUS To Protect	5-8
2. Configure the Switch To Access a RADIUS Server	5-10
3. Configure the Switch's Global RADIUS Parameters	5-12
Local Authentication Process	5-16
Controlling Web Browser Interface Access When Using RADIUS Authentication	5-17
Configuring RADIUS Accounting	5-17
Operating Rules for RADIUS Accounting	5-19
Steps for Configuring RADIUS Accounting	5-19
Viewing RADIUS Statistics	5-25
General RADIUS Statistics	5-25
RADIUS Authentication Statistics	5-27
RADIUS Accounting Statistics	5-28
Changing RADIUS-Server Access Order	5-30
Messages Related to RADIUS Operation	5-32

Overview

Feature	Default	Menu	CLI	Web
Configuring RADIUS Authentication	None	n/a	5-6	n/a
Configuring RADIUS Accounting	None	n/a	5-17	n/a
Viewing RADIUS Statistics	n/a	n/a	5-25	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 ProCurve 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 5-17.

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

Terminology

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

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

Host: See **RADIUS Server**.

NAS (Network Access Server): In this case, a ProCurve 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 ProCurve 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 5-25). 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 5-30.)
- 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 ProCurve 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.

```
ProCurve> show authentication
```

Status and Counters - Authentication Information

Login Attempts : 3
 Respect Privilege : Disabled

Access Task	Login Primary	Login Secondary	Enable Primary	Enable 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. ProCurve 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	5-8
< console telnet ssh web > < enable login > radius	5-8
< local none >	5-8
[no] radius-server host < <i>IP-address</i> >	5-10
[auth-port < <i>port-number</i> >]	5-10
[acct-port < <i>port-number</i> >]	5-10, 5-20
[key < <i>server-specific key-string</i> >]	5-10
[no] radius-server key < <i>global key-string</i> >	5-12
radius-server timeout < 1 - 15 >	5-12
radius-server retransmit < 1 - 5 >	5-12
[no] radius-server dead-time < 1 - 1440 >	5-14
show radius	5-25
[< host < <i>ip-address</i> >]	5-26
show authentication	5-27
show radius authentication	5-28

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:** The number of retries when there is no server response to a RADIUS authentication request. (Default: 3; range of 1 to 5.)
 - **Server Dead-Time:** The period during which the switch will not send new authentication requests to a RADIUS server that has failed to respond to a previous request. This avoids a wait for a request to time

out on a server that is unavailable. If you want to use this feature, select a dead-time period of 1 to 1440 minutes. (Default: 0—disabled; range: 1 - 1440 minutes.) If your first-choice server was initially unavailable, but then becomes available before the dead-time expires, you can nullify the dead-time by resetting it to zero and then trying to log on again. As an alternative, you can reboot the switch, (thus resetting the dead-time counter to assume the server is available) and then try to log on again.

- **Number of Login Attempts:** This is 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 5-17.)

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 Access Control (802.1X)” on page 8-1.

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 > < 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.)

[< local | none >]

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

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

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

Status and Counters - Authentication Information

Login Attempts : 3
Respect Privilege : Disabled

Access Task	Login Primary	Login Secondary	Enable Primary	Enable 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 5-2. 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 5-16.

2. Configure the Switch To Access a RADIUS Server

This section describes how to configure the switch to interact with a RADIUS server for both authentication and accounting services.

Note

If you want to configure RADIUS accounting on the switch, go to page 5-17: “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 5-30.)*

[auth-port < port-number >]

*Optional. Changes the UDP destination port for authentication requests to the specified RADIUS server (host). If you do not use this option with the **radius-server host** command, the switch automatically assigns the default authentication port number. The **auth-port** number must match its server counterpart. (Default: 1812)*

[acct-port < port-number >]

*Optional. Changes the UDP destination port for accounting requests to the specified RADIUS server. If you do not use this option with the **radius-server host** command, the switch automatically assigns the default accounting port number. The **acct-port** number must match its server counterpart. (Default: 1813)*

[key < key-string >]

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

no radius-server host < ip-address > key

*Use the **no** form of the command to remove the key for a specified server.*

For example, suppose you have configured the switch as shown in figure 5-3 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”.

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

  Server IP Addr  Auth  Acct  Encryption Key
  -----
  10.33.18.127   1812  1813  TempKey01
```

Figure 5-3. Sample Configuration for RADIUS Server Before Changing the Key and Adding Another Server

To make the changes listed prior to figure 5-3, you would do the following:

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

  Server IP Addr  Auth  Acct  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 5-4. 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 5-30.

3. Configure the Switch's Global RADIUS Parameters

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

- **Number of login attempts:** In a given session, specifies how many tries at entering the correct username and password pair are allowed before access is denied and the session terminated. (This is a general **aaa authentication** parameter and is not specific to RADIUS.)
- **Global server key:** The server key the switch will use for contacts with all RADIUS servers for which there is not a server-specific key configured by **radius-server host < ip-address > key < key-string >**. This key is optional if you configure a server-specific key for each RADIUS server entered in the switch. (Refer to “2. Configure the Switch To Access a RADIUS Server” on page 5-10.)
- **Server timeout:** Defines the time period in seconds for authentication attempts. If the timeout period expires before a response is received, the attempt fails.
- **Server dead time:** Specifies the time in minutes during which the switch avoids requesting authentication from a server that has not responded to previous requests.
- **Retransmit attempts:** If the first attempt to contact a RADIUS server fails, specifies how many retries you want the switch to attempt on that server.

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

Specifies how many tries for entering the correct 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)

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.

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

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

```

ProCurve# show authentication

Status and Counters - Authentication Information

(Login Attempts : 2) ←
Respect Privilege : Disabled

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

Access Task | Login      Login      Enable     Enable
              | 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

ProCurve# show radius

Status and Counters - General RADIUS Information

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

Global RADIUS parameters from figure 5-5.

Server IP Addr  Auth Port  Acct Port  Encryption Key
-----+-----
10.33.18.127    1812  1813  source0127
(10.33.18.119)  1812  1813
(10.33.18.151)  1812  1813 ←

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

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

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

Configuring RADIUS Accounting

RADIUS Accounting Commands	Page
[no] radius-server host < <i>ip-address</i> >	5-20
[acct-port < <i>port-number</i> >]	5-20
[key < <i>key-string</i> >]	5-20
[no] aaa accounting < exec network system > < start-stop stop-only > radius	5-23
[no] aaa accounting update periodic < 1 - 525600 > (<i>in minutes</i>)	5-24
[no] aaa accounting suppress null-username	5-24
show accounting	5-28
show accounting sessions	5-29
show radius accounting	5-29

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 5-5 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-Session-Id
- Acct-Status-Type
- Acct-Terminate-Cause
- Acct-Authentic
- Acct-Delay-Time
- Acct-Input-Packets
- Acct-Output-Packets
- Acct-Input-Octets
- Nas-Port
- Acct-Output-Octets
- Acct-Session-Time
- Username
- Service-Type
- NAS-IP-Address
- NAS-Identifier
- Called-Station-Id

(For 802.1X information for the switch, refer to “Configuring Port-Based Access Control (802.1X)” on page 8-1.)

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

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

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

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

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 5-30.)
- 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 “2. Configure the Switch To Access a RADIUS Server” on page 5-10.
- 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 5-10. (Default: null)
2. Configure accounting types and the controls for sending reports to the RADIUS server.
 - **Accounting types:** exec (page 5-18), network (page 5-18), or system (page 5-18)
 - **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 5-10. 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 5-10.)

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.).

```
ProCurve (config)# radius-server host 10.33.18.151 acct-port 1750 key source0151
ProCurve (config)# write mem
ProCurve (config)# show radius
Status and Counters - General RADIUS Information
  Deadtime (min) : 5
  Timeout (secs) : 3
  Retransmit Attempts : 2
  Global Encryption Key :

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

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

Figure 5-7. Example of Configuring for a RADIUS Server with a Non-Default Accounting UDP Port Number

The radius-server command as shown in figure 5-7, 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 Access Control (802.1X)” on page 8-1.

Determine how you want the switch to send accounting data to a RADIUS server:

■ **Start-Stop:**

- Send a start record accounting notice at the beginning of the accounting 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 5-22) 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 5-22) 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 > < 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:

```
ProCurve(config)# aaa accounting exec start-stop radius
ProCurve(config)# aaa accounting system stop-only radius
ProCurve(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
```

The diagram shows a terminal session with three main parts. The first part shows two configuration commands: 'aaa accounting exec start-stop radius' and 'aaa accounting system stop-only radius'. Callout boxes with arrows point to these lines, stating 'Configures exec and system accounting and controls.' The second part shows the 'show accounting' command output, which is enclosed in a dashed box. A callout box with an arrow points to this output, stating 'Summarizes the switch's accounting configuration.' The third part is a callout box with two arrows pointing to the 'Exec' and 'System' rows of the output table, stating 'Exec and System accounting are active. (Assumes the switch is configured to access a reachable'.

Figure 5-8. 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 5-8, suppose that you wanted the switch to:

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

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

ProCurve(config)# show accounting
Status and Counters - Accounting Information
Interval(min) : 10
Suppress Empty User : No

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

• Update Period
• Suppress Unknown User

Figure 5-9. 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 5-17.)*

```
ProCurve(config)# show radius
Status and Counters - General RADIUS Information
  Deadtme(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 5-10. Example of General RADIUS Information from Show Radius Command

```

ProCurve(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 5-11. RADIUS Server Information From the Show Radius Host Command

Table 5-2. Values for Show Radius Host Output (Figure 5-11)

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.

Term	Definition
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 5-17.)*

```

ProCurve> show authentication
Status and Counters - Authentication Information
Login Attempts : 2
-----+-----
Access Task | Login      Login      Enable     Enable
            | Primary    Secondary  Primary    Secondary
-----+-----
Console     | Local      None       Local      None
Telnet      | Radius     Local      Radius     Local
Port-Access | Local
SSH         | Radius     Local      Radius     Local
  
```

Figure 5-12. Example of Login Attempt and Primary/Secondary Authentication Information from the Show Authentication Command

```
ProCurve (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 5-13. 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.

```
ProCurve # 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 5-14. Listing the Accounting Configuration in the Switch


```
ProCurve# 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 5-15. Example of RADIUS Accounting Information for a Specific Server

```
ProCurve# show accounting sessions

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

Figure 5-16. 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:

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

  Server IP Addr  Auth Port  Acct 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 5-17. 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.

```
ProCurve(config)# no radius host 10.10.10.003
ProCurve(config)# no radius host 10.10.10.001
ProCurve(config)# radius host 10.10.10.003
ProCurve(config)# radius host 10.10.10.001
ProCurve(config)# show radius
```

Status and Counters - General RADIUS Information

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

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

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

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

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

Figure 5-18. Example of New RADIUS Server Search Order

Messages Related to RADIUS Operation

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

Configuring Secure Shell (SSH)

Contents

Overview	6-2
Terminology	6-4
Prerequisite for Using SSH	6-5
Public Key Formats	6-5
Steps for Configuring and Using SSH for Switch and Client Authentication	6-6
General Operating Rules and Notes	6-8
Configuring the Switch for SSH Operation	6-9
1. Assign Local Login (Operator) and Enable (Manager) Password .	6-9
2. Generate the Switch's Public and Private Key Pair	6-10
3. Provide the Switch's Public Key to Clients	6-13
4. Enable SSH on the Switch and Anticipate SSH Client Contact Behavior	6-15
5. Configure the Switch for SSH Authentication	6-18
6. Use an SSH Client To Access the Switch	6-22
Further Information on SSH Client Public-Key Authentication	6-22
Messages Related to SSH Operation	6-28

Overview

Feature	Default	Menu	CLI	Web
Generating a public/private key pair on the switch	No	n/a	page 6-10	n/a
Using the switch's public key	n/a	n/a	page 6-13	n/a
Enabling SSH	Disabled	n/a	page 6-15	n/a
Enabling client public-key authentication	Disabled	n/a	pages 6-19, 6-22	n/a
Enabling user authentication	Disabled	n/a	page 6-18	n/a

The ProCurve 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.)

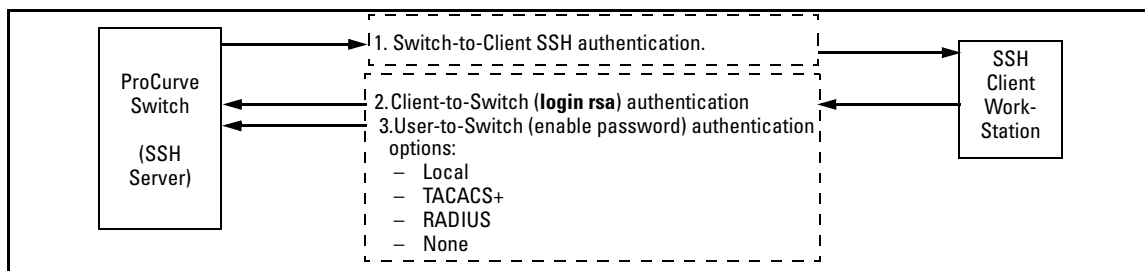


Figure 6-1. Client Public Key Authentication Model

Note

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

Switch SSH and User Password Authentication . This option is a subset of the client public-key authentication show in figure 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

SSH on the ProCurve switches covered in this guide supports these data encryption methods:

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

Note

The ProCurve switches covered in this guide use the RSA algorithm for internally generated keys (v2 shared host key). However, ProCurve switches support both RSA and DSA/DSS keys for client authentication. All references to either a public or private key mean keys generated using these algorithms unless otherwise noted

Terminology

- **SSH Server:** A ProCurve switch with SSH enabled.
- **Key Pair:** A pair of keys generated by the switch or an SSH client application. Each pair includes a public key, that can be read by anyone and a private key, that is held internally in the switch or by a client.
- **PEM (Privacy Enhanced Mode):** Refers to an ASCII-formatted client public-key that has been encoded for portability and efficiency. SSHv2 client public-keys are typically stored in the PEM format. See figure 6-3 for an example of a PEM-encoded ASCII key.
- **Private Key:** An internally generated key used in the authentication process. A private key generated by the switch is not accessible for viewing or copying. A private key generated by an SSH client application is typically stored in a file on the client device and, together with its public key counterpart, can be copied and stored on multiple devices.
- **Public Key:** An internally generated counterpart to a private key. A device's public key is used to authenticate the device to other devices.
- **Enable Level:** Manager privileges on the switch.
- **Login Level:** Operator privileges on the switch.
- **Local password or username:** A Manager-level or Operator-level password configured in the switch.
- **SSH Enabled:** (1) A public/private key pair has been generated on the switch (**crypto key generate ssh [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. Generate the Switch's Public and Private Key Pair” on page 6-10 and “4. Enable SSH on the Switch and Anticipate SSH Client Contact Behavior” on page 6-15.)

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 6-2), 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 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"A1B3Hz1y2+orEhYL . . . Q8D8qDH1ozu1c="*** 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 ten client keys.) This topic is covered under “To Create a Client-Public-Key Text File” on page 6-24.

B. Switch Preparation

1. Assign a login (Operator) and enable (Manager) password on the switch (page 6-9).
2. Generate a public/private key pair on the switch (page 6-10).

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 6-13).
4. Enable SSH on the switch (page 6-15).
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 authentication method is local, the secondary method must be none.

- Option B:

- Primary: Client public-key authentication (**login public-key** — page 6-22)
 - 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 authentication method is local, the secondary method 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 ten 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.
- On ProCurve switches that support stacking, when stacking is enabled, SSH provides security only between an SSH client and the stack manager. Communications between the stack commander and stack members are not secure.
- The switch does not support outbound SSH sessions. Thus, if you Telnet from an SSH-secure switch to another SSH-secure switch, *the session is not secure*.

Configuring the Switch for SSH Operation

SSH-Related Commands in This Section	Page
show ip ssh	6-17
show crypto client-public-key [manager operator] [keylist-str] [< babble fingerprint >]	6-25
show crypto host-public-key [< babble fingerprint >]	6-14
show authentication	6-21
crypto key < generate zeroize > ssh [rsa]	6-11
ip ssh	6-16
port < 1 - 65535 default >	6-16
timeout < 5 - 120 >	6-16
aaa authentication ssh	
login < local tacacs radius public-key >	6-18, 6-20
< local none >	6-18
enable < tacacs radius local >	6-18
< local none >	6-18
copy tftp pub-key-file <tftp server IP> <public key file> [<append manager operator>]	6-25
clear crypto client-public-key [keylist-str]	6-27

1. Assign Local Login (Operator) and Enable (Manager) Password

At a minimum, ProCurve 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 >

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

Figure 6-4. Example of Configuring Local Passwords

2. Generate 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.)

Notes

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

Removing (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 RSA Host Key Pair.

Because the host key pair is stored in flash instead of the running-config file, it is not necessary to use **write memory** to save the key pair. Erasing the key pair automatically disables SSH.

Syntax: `crypto key generate ssh [rsa]`

Generates a public/private key pair for the switch. If a switch key pair already exists, replaces it with a new key pair. (See the Note, above.)

`crypto key zeroize ssh [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 6-14.)

`[fingerprint]`

Displays fingerprints of the switch's public key in hexadecimal format. (See "Displaying the Public Key" on page 6-14.)

For example, to generate and display a new key:

```
ProCurve(config)# crypto key generate ssh rsa
Installing new RSA key.  If the key/entropy cache is
depleted, this could take up to a minute.
ProCurve(config)# show crypto host-public-key

-----
SSH host public key file
Version 1 format:

|896 35 3219295003103011452137203169501232714847265325085720757925409572738582167|
|49173126937413223781326827636154399173519641900117298772018339016754333892248319|
|41759125186557710233731689070801858880718460531164552600040416069890120011153581|
|9449254242176260739141950918771764467
Version 2 format:

|ssh-rsa AAAAB3NzaClyc2EAAAABTwAAAHEAnAAApdhq13Jynrs7j4lDUm8ivVm8ld2mZU5e+YZWp/T6|
|QzP2RsDDMZLbAHHIBrxPLjW/bRogpYD0lWuVOhTojEVjqeVuXbwmdDnyOgBc06olePwdrbQ+FZevERiA|
|JYG3C8NCzCRD/djXeI7FmRps8w==
-----
```

Figure 6-5. Example of Generating a Public/Private Host Key Pair for the Switch

The 'show crypto host-public-key' displays data in two different formats because your client may store it in either of these formats after learning the key. If you wish to compare the switch key to the key as stored in your client's known-hosts file, note that the formatting and comments need not match. For version 1 keys, the three numeric values bit size, exponent <e>, and modulus <n> must match; for PEM keys, only the PEM-encoded string itself must 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.

3. Provide 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 generated by the switch consists of three parts, separated by one blank space each:

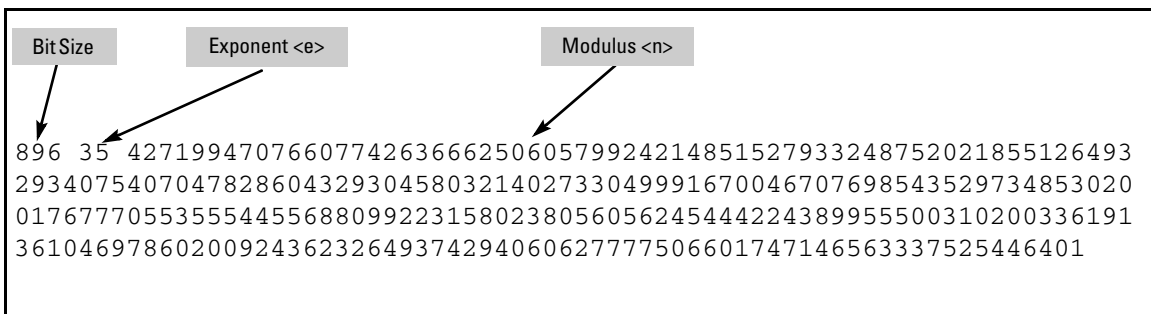


Figure 6-6. Example of a Public Key Generated by the Switch

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

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

1. Use a terminal application such as HyperTerminal to display the switch's public key with the **show 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 in 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:

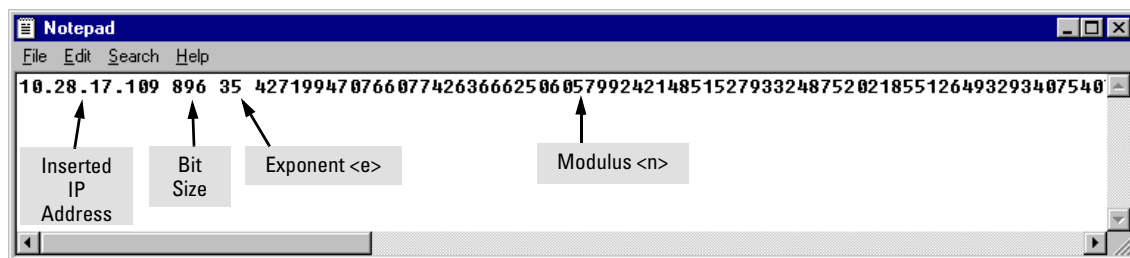


Figure 6-8. Example of a Switch Public Key Edited To Include the Switch's IP Address

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 6-13.)
- **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:

```
ProCurve(config)# show crypto host-public-key babble
896 xozik-kobaf-daroh-fygas-byveb-bymiz-nupap-povaz-cesin-kafec-rixux
   host_sshl
896 xefes-hikot-kyher-cukuz-balah-gezos-gumym-rezif-horib-cicyp-poxyx
   host_ssh2.pub
ProCurve(config)# show crypto host-public-key fingerprint
896 53:c0:14:75:72:84:90:cc:c8:ba:5e:ca:92:fc:c7:5c host_sshl
896 bf:fb:8a:d0:10:5a:48:57:61:f9:8a:6a:61:13:8a:fb host_ssh2.pub
```



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. Enable SSH on the Switch and Anticipate SSH Client Contact Behavior

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

Note

Before enabling SSH on the switch you must generate the switch’s public/private key pair. If you have not already done so, refer to “2. Generate the Switch’s Public and Private Key Pair” on page 6-10.

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. Configure the Switch for SSH Authentication” on page 6-18.

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

Note

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

To enable SSH on the switch.

1. Generate a public/private key pair if you have not already done so. (Refer to "2. Generate the Switch's Public and Private Key Pair" on page 6-10.)
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 6-11).

Syntax: [no] ip ssh

Enables or disables SSH on the switch.

[port < 1-65535 | default >]

*The TCP port number for SSH connections (default: 22). **Important:** See "Note on Port Number" on page 6-17.*

[timeout < 5 - 120 >]

The SSH login timeout value (default: 120 seconds).

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

Note on Port Number

ProCurve 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 ProCurve switches are 49, 80, 1506, and 1513.

```
ProCurve(config)# ip ssh
ProCurve(config)# show ip ssh
```

SSH Enabled	: Yes
SSH Version	: 1-or-2
IP Port Number	: 22
Timeout (sec)	: 120
Server Key Size (bits)	: 512

```
Ses Type      | Protocol  Source IP and Port
-----+-----
1  console   |
2  telnet    |
3  ssh       | SSH v2    12.255.255.255:18
4  inactive  |
```

Enables SSH on the switch.

Lists the current SSH configuration and status.

The switch uses these five settings internally for transactions with clients. See the **Caution** on page 6-18.

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

Figure 6-10. Example of Enabling IP SSH and Listing the SSH Configuration and Status

Caution

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

SSH does not protect the switch from unauthorized access via the web interface, Telnet, SNMP, or the serial port. While web and Telnet access can be restricted by the use of passwords local to the switch, if you are unsure of the security this provides, you may want to disable web-based and/or Telnet access (**no web-management** and **no telnet**). If you need to increase SNMP security, you should use SNMP version 3 only. If you need to increase the security of your web interface, refer to chapter 7, “Configuring Secure Socket Layer (SSL)”. Another security measure is to use the Authorized IP Managers feature described in the switch’s *Management and Configuration Guide*. To protect against unauthorized access to the serial port (and the Clear button, which removes local password protection), keep physical access to the switch restricted to authorized personnel.

5. Configure 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 (page 6-19) 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 6-22

Note

ProCurve 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. Assign Local Login (Operator) and Enable (Manager) Password” on page 6-9.*

Option A: Configuring SSH Access for Password-Only SSH

Authentication. When configured with this option, the switch uses its public key to authenticate itself to a client, but uses only passwords for client authentication.

Syntax: `aaa authentication ssh login < local | tacacs | radius >[< local | none >]`

*Configures a password method for the primary and secondary login (Operator) access. If you do not specify an optional secondary method, it defaults to **none**.*

*Note: If the primary access is **local**, the secondary access cannot be **local**.*

`aaa authentication ssh enable < local | tacacs | radius>[< local | none >]`

*Configures a password method for the primary and secondary enable (Manager) access. If you do not specify an optional secondary method, it defaults to **none**.*

*Note: If the primary access is **local**, the secondary access cannot be **local**.*

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 6-22.)

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

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

Copies a public key file into the switch.

aaa authentication ssh login public-key

*Configures the switch to authenticate a client public-key at the login level with an optional secondary password method (Default: **none**).*

*Note: The secondary access cannot be **local**.*

Caution

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

Syntax: aaa authentication ssh enable < local | tacacs | radius > < local | none >

*Configures a password method for the primary and secondary enable (Manager) access. If you do not specify an optional secondary method, it defaults to **none**.*

*Note: If the primary access is **local**, the secondary access cannot be **local**.*

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


```

ProCurve(config)# password manager user-name leader
New password for Manager: *****
Please retype new password for Manager: *****
ProCurve(config)# [aaa authentication ssh login public-key none]
ProCurve(config)# [aaa authentication ssh enable tacacs local]
ProCurve(config)# [copy tftp pub-key-file 12.255.255.255 Client-key.pub]
ProCurve(config)# write memory
  
```

Configures Manager user-name and password.

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

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

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

Figure 6-11. Configuring for SSH Access Requiring a Client Public-Key Match and Manager Passwords

Figure 6-12 shows how to check the results of the above commands.

```

ProCurve (config)# show authentication
  
```

Status and Counters - Authentication Information

Login Attempts : 3

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

Lists the current SSH authentication configuration.

Shows the contents of the public key file downloaded with the copy tftp command in figure 6-11. In this example, the file contains two client public-keys.

Client Key Index Number

```

ProCurve (config)# show crypto client-public-key
[ 0] "Maden name [1024-bit rsa, Local_crypto @ Localcrypto, Thu Nov 07 2002 21:25:42]" ssh-rsa AAAAB3NzaClyc2EAAAADAQABAAQgQCz9oNfQxMHUFEC6frS1Sa4Uh1EFznFhQcmgP29HXYP6NR/1Q0UmActrFU+QD11EtM/YM9FrN/XvZH/kIxTdEc5exFX/S10tcRaFYzI9UjK80dBmqvBGKB
IyVEbCVw1qdAqbkaEX3d/WaPS2xArLCFHsTZhmCvQTZDOGAB1frlcw==
[ 1] "[768-bit rsa, Local_crypto @ Localcrypto, Mon Dec 16 2002 23:01:51]" ssh-rsa
AAAAB3NzaClyc2EAAAADAQABAAQgQD0tmzA32JBgeuFJNOiXI3bfooPKZ09JKCPQcXEVk7N+eKf9MOX
vnmfFuEpw/fpqhlvsE66n8FDu7W/B2tKH/tgQLFqX7GiVcxNGhLiNO/pg5AuEym8Enc1Gu/LqAM9daM=
  
```

Figure 6-12. SSH Configuration and Client-Public-Key Listing From Figure 6-11

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. Configure the Switch for SSH Authentication" on page 6-18 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 RSA or DSA 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. Configure the Switch for SSH Authentication" on page 6-18.)

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

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

Using client public-key authentication requires these steps:

1. Generate a public/private key pair for each client you want to have SSH access to the switch. This can be a separate key for each client or the same key copied to several clients.
2. Copy the public key for each client into a client-public-key text file.
3. Use **copy tftp** to copy the client-public-key file into the switch. Note that 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.

Configuring Secure Shell (SSH)

Further Information on SSH Client Public-Key Authentication

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

Bit Size	Exponent <e>	Modulus <n>	Comment
1024	35	1140740666170144690796380365284018053912704374511148288250928555011016860308260389591468963065690359820412220255425432827643299433440329635043810210989476474605645572227682031607648603664020534703408371002884293231503492265409355321119922465153140745413543765609589968291386053556814705585051025488575846923	smith@support.cairns.com

Figure 6-13. Example of a Client Public Key

Notes

Comments in public key files, such as **smith@support.cairns.com** in figure 6-13, 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 6-13 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 6-13. 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 only	
Maximum Supported Public Key Length	3072 bits	Shorter key lengths allow faster operation, but also mean diminished security.
Maximum Key Size	1024 characters	Includes the bit size, public index, modulus, any comments, <CR>, <LF>, and all blank spaces. If necessary, you can use an editor application to verify the size of a key. For example, placing a client-public-key into a Word for Windows text file and clicking on File Properties 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 authentication, copy a public key for each of these clients (up to ten) into the file. Each key should be separated from the preceding key by a <CR><LF>.
3. Copy the client-public-key file into a TFTP server accessible to the switch.

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

- One or more client-generated public keys. 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 individual 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@fellow** at the end of the key in figure 6-13 on page 6-24.)

Configuring Secure Shell (SSH)

Further Information on SSH Client Public-Key Authentication

Syntax: copy tftp pub-key-file <ip-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 adds 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 **manager** option selects the manager public keys.*

*The **operator** option selects operator public keys.*

*The **keylist-str** option allows you to select keys to display (a comma-delimited list).*

*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.*

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:

```
ProCurve (config)# copy tftp pub-key-file 10.38.252.195 Clientkeys.txt
ProCurve (config)# show crypto client-public-key
0."Maden name [1024-bit rsa, Jamie_wilson@Jamiewilson, Thu Nov 07 2002 21:25:4
2]" ssh-rsa AAAAB3NzaClyc2EAAAADAQABAAQgQCz9cNfqxMHUFEC6frSulSa4UhlEFznFhQqmgP2
9HXVp6NR/lQOUmACtrFU+QD11EtM/YM9FrN/XvZH/kIxTdEc5exFX/S10tcRaFYzI9UjK80dBmQvBGKB
IyVEbCVwlqdaqbkaEX3d/WaPS2xArLCFHsTZhnCvQTZD0GABlfrlcw==
1."[768-bit rsa, Jamie_wilson@Jamiewilson, Mon Dec 16 2002 23:01:51]" ssh-rsa
AAAAB3NzaClyc2EAAAADAQABAAQgYQD0tmzA32JBgeuFJN0iXI3bfooPKZ09JKCPQcXEVk7N+eKf9MOX
vnmfFuEpw/fpqhlvsE66n8FDu7W/B2tKH/tqQLFcx7GiVcxNGhLiN0/pq5AuEym8Enc1Gu/LgAM9daM=
```

Key Index Number

Figure 6-14. Example of Copying and Displaying a Client Public-Key File Containing Two Client Public Keys

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

- You can 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.

Caution

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

Messages Related to SSH Operation

Message	Meaning
00000K Peer unreachable.	Indicates an error in communicating with the tftp server or not finding the file to download. Causes include such factors as: <ul style="list-style-type: none">• Incorrect IP configuration on the switch• Incorrect IP address in the command• Case (upper/lower) error in the filename used in the command• Incorrect configuration on the TFTP server• The file is not in the expected location.• Network misconfiguration• No cable connection to the network
00000K Transport error.	Indicates the switch experienced a problem when trying to copy tftp the requested file. The file may not be in the expected directory, the filename may be misspelled in the command, or the file permissions may be wrong.
Cannot bind reserved TCP port <port-number>.	The ip ssh port command has attempted to configure a reserved TCP port. Use the default or select another port number. See "Note on Port Number" on page 6-17.

Message	Meaning
Client public key file corrupt or not found. Use 'copy tftp pub-key-file <ip-addr> <filename>' to download new file.	The client key does not exist in the switch. Use copy tftp to download the key from a TFTP server.
Download failed: overlength key in key file.	The public key file you are trying to download has one of the following problems: <ul style="list-style-type: none">• A key in the file is too long. The maximum key length is 1024 characters, including spaces. This could also mean that two or more keys are merged together instead of being separated by a <CR><LF>.• There are more than ten public keys in the key file and switch total. Delete some keys from the switch or file. The switch does not detect duplicate keys.• One or more keys in the file is corrupted or is not a valid rsa public key. Refer to “To Create a Client-Public-Key Text File” on page 24 for information on client-public-key properties.
Download failed: too many keys in key file.	
Download failed: one or more keys is not a valid public key.	
Error: Requested keyfile does not exist.	The client key does not exist in the switch. Use copy tftp to download the key from a TFTP server.
Generating new RSA host key. If the cache is depleted, this could take up to two minutes.	After you execute the crypto key generate ssh [rsa] command, the switch displays this message while it is generating the key.
Host RSA key file corrupt or not found. Use 'crypto key generate ssh rsa' to create new host key.	The switch's key is missing or corrupt. Use the crypto key generate ssh [rsa] command to generate a new key for the switch.

—This page is intentionally unused—

Configuring Secure Socket Layer (SSL)

Contents

Overview	7-2
Terminology	7-3
Prerequisite for Using SSL	7-5
Steps for Configuring and Using SSL for Switch and Client Authentication	7-5
General Operating Rules and Notes	7-6
1. Assign Local Login (Operator) and Enable (Manager) Password .	7-7
2. Generate the Switch's Server Host Certificate	7-9
3. Enable SSL on the Switch and Anticipate SSL Browser Contact Behavior	7-17
Common Errors in SSL Setup	7-21

Overview

Feature	Default	Menu	CLI	Web
Generating a Self Signed Certificate on the switch	No	n/a	page 7-9	page 7-13
Generating a Certificate Request on the switch	No	n/a	n/a	page 7-15
Enabling SSL	Disabled	n/a	page 7-17	page 7-19

The ProCurve switches covered by this manual 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.

Note

ProCurve 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.

Note

SSL in ProCurve switches is based on the OpenSSL software toolkit. For more information on OpenSSL, visit <http://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.

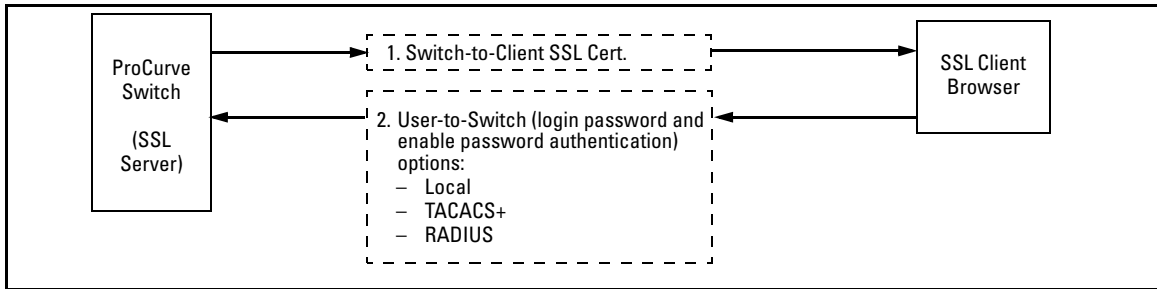


Figure 7-1. Switch/User Authentication

SSL on the ProCurve switches supports these data encryption methods:

- 3DES (168-bit, 112 Effective)
- DES (56-bit)
- RC4 (40-bit, 128-bit)

Note:

ProCurve switches use RSA public key algorithms and Diffie-Hellman. All references to a key mean keys generated using these algorithms unless otherwise noted

Terminology

- **SSL Server:** A ProCurve 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 Procurve 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.)

Note:

The latest versions of Microsoft Internet Explorer and Netscape web browser support SSL and TLS functionality. See the browser documentation for additional details

B. Switch Preparation

1. Assign a login (Operator) and enable (Manager) password on the switch (page 7-7).
2. Generate a host certificate on the switch (page 7-9).
 - 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 7-17).
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
- On ProCurve switches that support stacking, when stacking is enabled, SSL provides security only between an SSL client and the stack manager. Communications between the stack commander and stack members is not secure.

Configuring the Switch for SSL Operation

SSL-Related CLI Commands in This Section	Page
web-management ssl	page 7-19
show config	page 7-19
show crypto host-cert	page 7-12
crypto key	
generate cert [rsa] <512 768 1024>	page 7-10
zeroize cert	page 7-10
crypto host-cert	
generate self-signed [arg-list]	page 7-10
zeroize	page 7-10

1. Assign Local Login (Operator) and Enable (Manager) Password

At a minimum, ProCurve 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 *Management and Configuration 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. Generate 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 their 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 certificates 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

Note:

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] < 512 | 768 | 1024 >`

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 on page 7-9.)

`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 512.

Note:

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.

Note:

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:

```

ProCurve (config)#crypto key generate cert 512
Installing new RSA key.  If the key/entropy cache is
depleted, this could take up to a minute.
ProCurve (config)#crypto host-cert generate self-signed
Validity start date [01/01/1970]: 01/01/2003
Validity end date   [01/01/2004]: 01/01/2004
Common name        [0.0.0.0]: 10.255.255.255
Organizational unit [Dept Name]: ProCurve Networking
Organization       [Company Name]: Hewlett-Packard
City or location    [City]: Roseville
State name          [State]: CA
Country code        [US]: US
  
```

Figure 7-3. Example of Generating a Self-Signed Server Host certificate on the CLI for the Switch.

Notes

“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:

```
ProCurve(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=ProCurve Network
Validity
  Not Before: Jan 1 00:00:00 2002 GMT
  Not After : Jan 1 23:59:59 2004 GMT
Subject: CN=10.255.255.255, L=Roseville, ST=Ca, C=US, O=Hewlett Packard, OU=ProCurve Network
Subject Public Key Info:
  Public Key Algorithm: rsaEncryption
  RSA Public Key: (512 bit)
  Modulus (512 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: md5WithRSAEncryption
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 *Management and Configuration Guide* for your switch.

To generate a self signed host certificate from the web browser interface:

- i. Select the **Security** tab then the **[SSL]** button. The SSL configuration screen is divided into two halves. The left half is used for 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 **Create Certificate/Certificate Request** radio button.
- iii. Select **Self-Signed** in the **Certificate Type** drop-down list.
- iv. Select the **RSA Key Size** desired. If you want to re-use the current certificate key, select **Current** from this list.
- v. Fill in the remaining certificate arguments. (Refer to “Comments on Certificate Fields.” on page 7-11.)
- vi. Click on the **[Apply Changes]** button to generate new certificate and key, if selected.

Note

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.

Configuring Secure Socket Layer (SSL) General Operating Rules and Notes

For example, to generate a new host certificate via the web browsers interface:

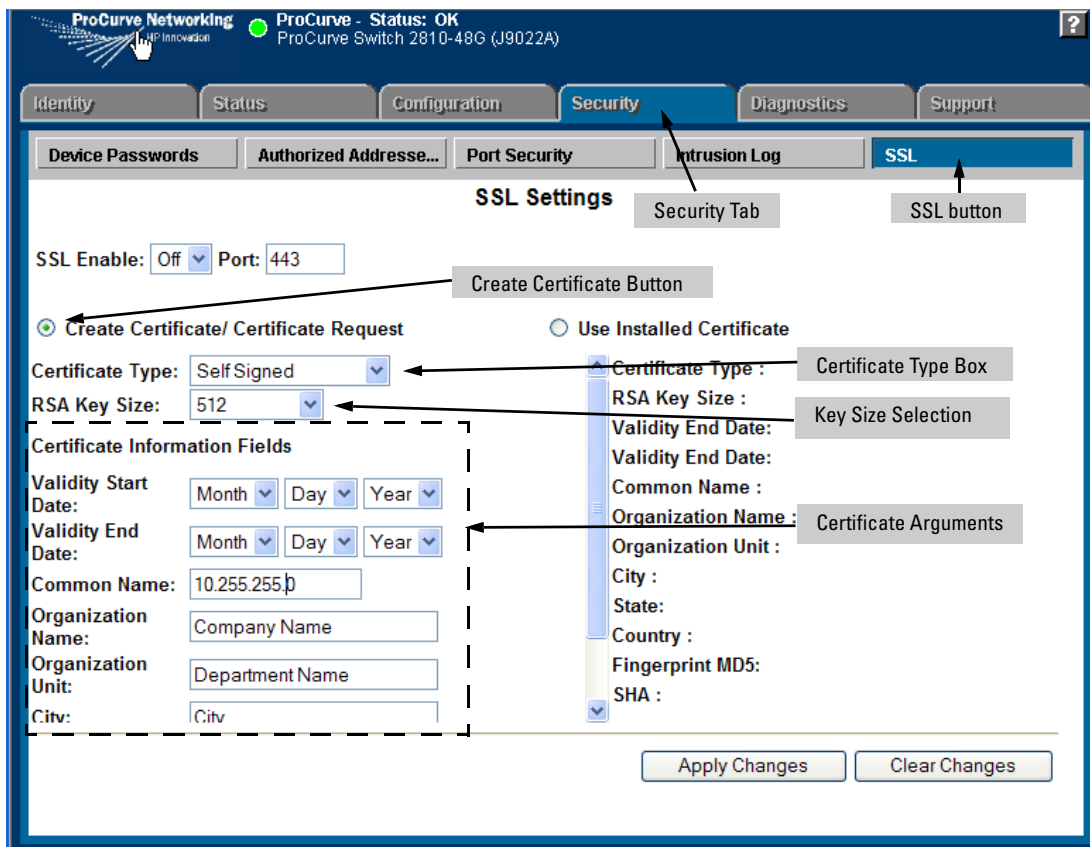


Figure 7-5. Self-Signed Certificate generation via SSL Web Browser Interface Screen

To view the current host certificate in the web browser interface:

1. Select the **Security** tab.
2. Select the **[SSL]** button.

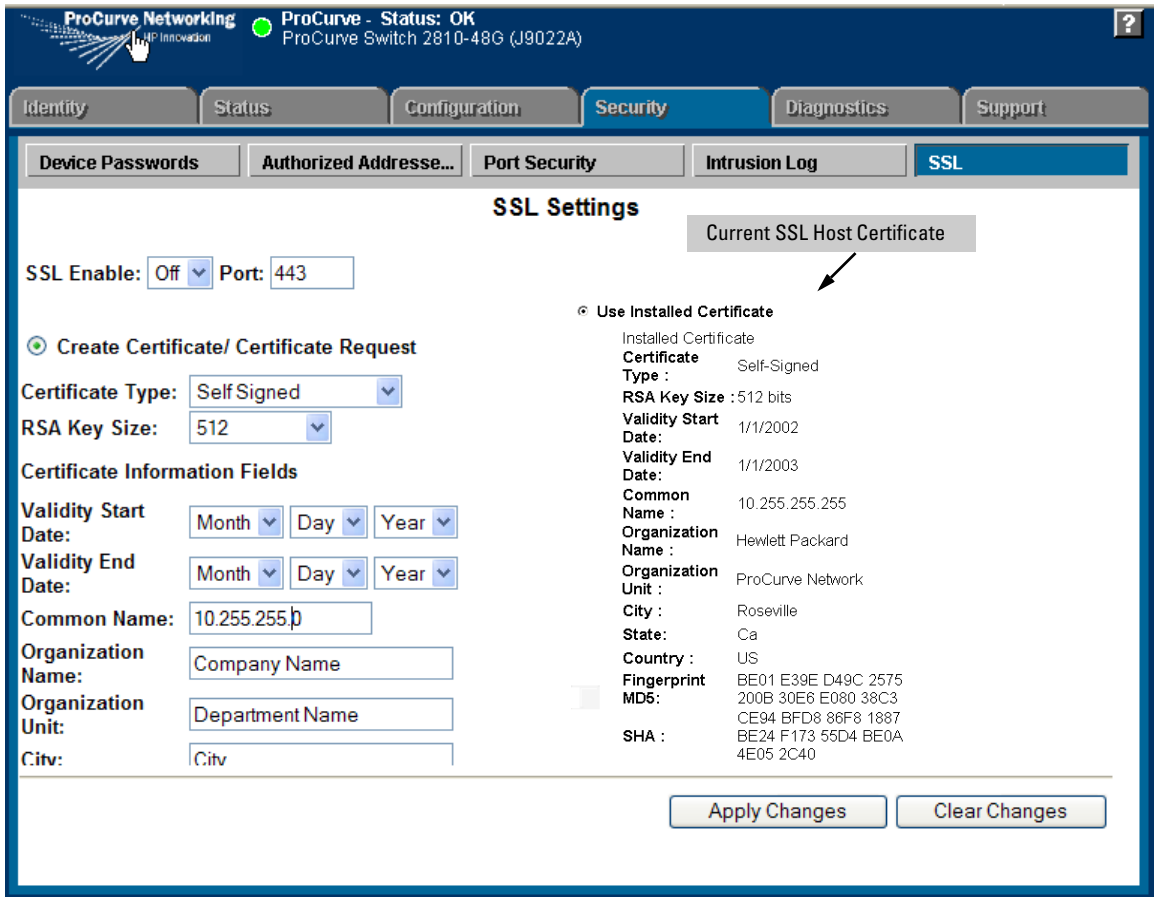


Figure 7-6. Web browser Interface showing current SSL Host Certificate

Generate a CA-Signed server host certificate with the Web Browser Interface

This section describes how 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 *Management and Configuration 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 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 want to re-use the current certificate key, select **Current** from this list.
- v. Fill in the remaining certificate arguments. (Refer to “Comments on Certificate Fields.” on page 7-11.)
- 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 none-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 the certificate into the lower text box.
- viii. Click on the **[Apply Changes]** button to install the certificate.

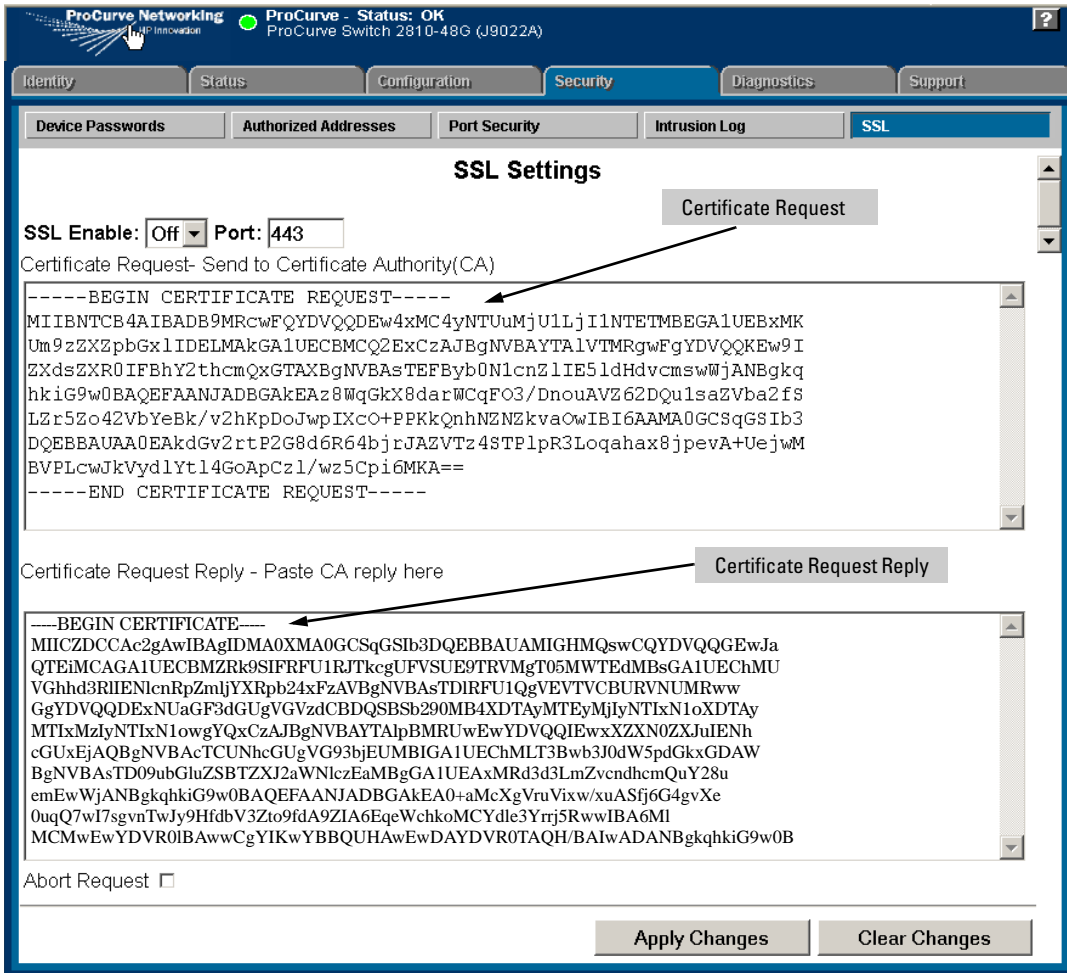


Figure 7-7. Example of a Certificate Request and Reply

3. Enable SSL on the Switch and Anticipate 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. The **no web-management ssl** command is used to disable SSL on the switch.

Note

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. Generate the Switch's Server Host Certificate" on page 7-9.

When configured for SSL, the switch uses its host certificate to authenticate itself to SSL clients, however unless you disable the standard 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.

Note

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 7-20.*

show config

*Shows status of the SSL server. When enabled, **web-management ssl** appears 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. Generate the Switch’s Server Host Certificate” on page 7-9.)
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 7-10).

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.

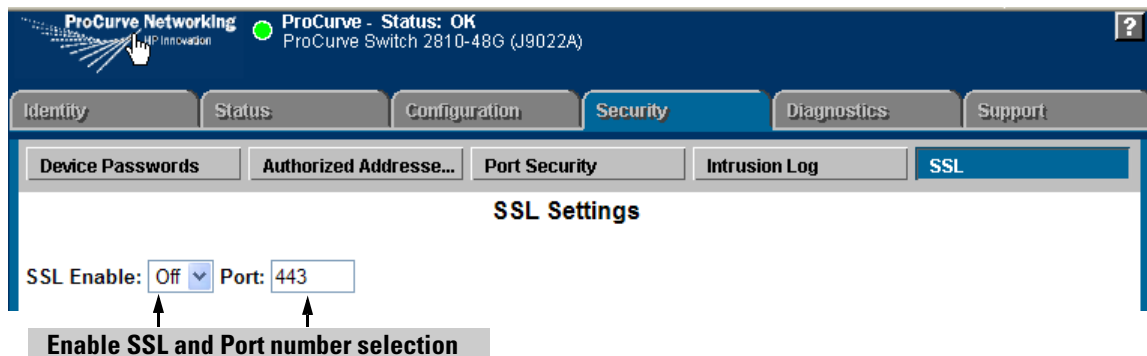


Figure 7-8. Using the web browser interface to enable SSL and select TCP port number

Note on Port Number

ProCurve 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 switch 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 7-10.)
Enabling SSL on the CLI or Web browser interface	<p>You have not generated a host certificate. (Refer to "Generate a Self-Signed Host Certificate with the Web browser interface" on page 7-13.)</p> <p>You may be using a reserved TCP port. (Refer to "Note on Port Number" on page 7-20.)</p>
Unable to Connect with SSL	<p>You may not have SSL enabled (Refer to "3. Enable SSL on the Switch and Anticipate SSL Browser Contact Behavior" on page 7-17.)</p> <p>Your browser may not support SSLv3 or TLSv1 or it may be disabled. (Refer to the documentation provided for your browser.)</p>

— This page is intentionally unused. —

Configuring Port-Based Access Control (802.1X)

Contents

Overview	8-3
Why Use Port-Based Access Control?	8-3
General Features	8-3
How 802.1X Operates	8-6
Authenticator Operation	8-6
Switch-Port Supplicant Operation	8-7
Terminology	8-8
General Operating Rules and Notes	8-10
General Setup Procedure for Port-Based Access Control (802.1X)	8-12
Do These Steps Before You Configure 802.1X Operation	8-12
Overview: Configuring 802.1X Authentication on the Switch	8-13
Configuring Switch Ports as 802.1X Authenticators	8-15
1. Enable 802.1X Authentication on Selected Ports	8-15
3. Configure the 802.1X Authentication Method	8-19
4. Enter the RADIUS Host IP Address(es)	8-20
5. Enable 802.1X Authentication on the Switch	8-20
802.1X Open VLAN Mode	8-21
Introduction	8-21
Use Models for 802.1X Open VLAN Modes	8-22
Operating Rules for Authorized-Client and Unauthorized-Client VLANs	8-25
Setting Up and Configuring 802.1X Open VLAN Mode	8-27
802.1X Open VLAN Operating Notes	8-31
Option For Authenticator Ports: Configure Port-Security To Allow Only 802.1X Devices	8-32

Configuring Port-Based Access Control (802.1X)

Contents

Configuring Switch Ports To Operate As Suplicants for 802.1X Connections to Other Switches	8-34
Displaying 802.1X Configuration, Statistics, and Counters	8-38
Show Commands for Port-Access Authenticator	8-38
Viewing 802.1X Open VLAN Mode Status	8-40
Show Commands for Port-Access Supplicant	8-43
How RADIUS/802.1X Authentication Affects VLAN Operation	8-44
Messages Related to 802.1X Operation	8-48

Overview

Feature	Default	Menu	CLI	Web
Configuring Switch Ports as 802.1X Authenticators	Disabled	n/a	page 8-15	n/a
Configuring 802.1X Open VLAN Mode	Disabled	n/a	page 8-21	n/a
Configuring Switch Ports to Operate as 802.1X Supplicants	Disabled	n/a	page 8-34	n/a
Displaying 802.1X Configuration, Statistics, and Counters	n/a	n/a	page 8-38	n/a
How 802.1X Affects VLAN Operation	n/a	n/a	page 8-44	n/a
RADIUS Authentication and Accounting	Refer to “RADIUS Authentication and Accounting” on page 5-1			

Why Use Port-Based Access Control?

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

General Features

802.1X on the ProCurve 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 or CHAP protocol.
 - Provision for enabling clients that do not have 802.1 supplicant software to use the switch as a path for downloading the software and initiating the authentication process (802.1X Open VLAN mode).
 - Supplicant implementation using CHAP authentication and independent username and password configuration on each port.
- Prevention of traffic flow in either direction on unauthorized ports.

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

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

Providing a Path for Downloading 802.1X Supplicant Software. For clients that do not have the necessary 802.1X supplicant software, there is also the option to configure the 802.1X Open VLAN mode. This mode allows you to assign such clients to an isolated VLAN through which you can provide the necessary supplicant software these clients need to begin the authentication process. (Refer to "802.1X Open VLAN Mode" on page 8-21.)

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 8-1. Example of an 802.1X Application

Accounting . The switch also provides RADIUS Network accounting for 802.1X access. Refer to “RADIUS Authentication and Accounting” on page 5-1.

How 802.1X Operates

Authenticator Operation

This operation provides security on a direct, point-to-point link between a single client and the switch, where both devices are 802.1X-aware. (If you expect desirable clients that do not have the necessary 802.1X supplicant software, you can provide a path for downloading such software by using the 802.1X Open VLAN mode—refer to “802.1X Open VLAN Mode” on page 8-21.) For example, suppose that you have configured a port on the switch for 802.1X authentication operation. If you then connect an 802.1X-aware client (supplicant) to the port and attempt to log on:

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

Switch-Port Supplicant Operation

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

- Switch “A” has port A1 configured for 802.1X supplicant operation.
- You want to connect port A1 on switch “A” to port B5 on switch “B”.




Figure 8-2. Example of Supplicant Operation

1. When port A1 on switch “A” is first connected to a port on switch “B”, or if the ports are already connected and either switch reboots, port A1 begins sending start packets to port B5 on switch “B”.
 - If, after the supplicant port sends the configured number of start packets, it does not receive a response, it assumes that switch “B” is not 802.1X-aware, and transitions to the authenticated state. If switch “B” is operating properly and is not 802.1X-aware, then the link should begin functioning normally, but without 802.1X security.
 - If, after sending one or more start packets, port A1 receives a request packet from port B5, then switch “B” is operating as an 802.1X authenticator. The supplicant port then sends a response/ID packet. Switch “B” forwards this request to a RADIUS server.
2. The RADIUS server then responds with an MD5 access challenge that switch “B” forwards to port A1 on switch “A”.
3. Port A1 replies with an MD5 hash response based on its username and password or other unique credentials. Switch “B” forwards this response to the RADIUS server.
4. The RADIUS server then analyzes the response and sends either a “success” or “failure” packet back through switch “B” to port A1.
 - A “success” response unblocks port B5 to normal traffic from port A1.

- A “failure” response continues the block on port B5 and causes port A1 to wait for the “held-time” period before trying again to achieve authentication through port B5.

Note

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

Terminology

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

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

Authentication Server: The entity providing an authentication service to the switch when the switch is configured to operate as an authenticator. In the case of an ProCurve 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 ProCurve switch applications, a device such as a switch that requires a supplicant to provide the proper credentials (username and password) before being allowed access to the network.

CHAP (MD5): Challenge Handshake Authentication Protocol.

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

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

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

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

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

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

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

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

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

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

member of that VLAN as long as at least one other port on the switch is statically configured as a tagged or untagged member of the same Unauthorized-Client VLAN.

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

General Operating Rules and Notes

- When a port on the switch is configured as either an authenticator or supplicant and is connected to another device, rebooting the switch causes a re-authentication of the link.
- When a port on the switch is configured as an authenticator, it will block access to a client that either does not provide the proper authentication credentials or is not 802.1X-aware. (You can use the optional 802.1X Open VLAN mode to open a path for downloading 802.1X supplicant software to a client, which enables the client to initiate the authentication procedure. Refer to “802.1X Open VLAN Mode” on page 8-21.)
- If a port on switch “A” is configured as an 802.1X supplicant and is connected to a port on another switch, “B”, that is not 802.1X-aware, access to switch “B” will occur without 802.1X security protection.
- You can configure a port as both an 802.1X authenticator *and* an 802.1X supplicant.
- If a port on switch “A” is configured as both an 802.1X authenticator *and* supplicant and is connected to a port on another switch, “B”, that is not 802.1X-aware, access to switch “B” will occur without 802.1X security protection, but switch “B” will not be allowed access to switch “A”. This means that traffic on this link between the two switches will flow from “A” to “B”, but not the reverse.

- If a client already has access to a switch port when you configure the port for 802.1X authenticator operation, the port will block the client from further network access until it can be authenticated.
- On a port configured for 802.1X with RADIUS authentication, if the RADIUS server specifies a VLAN for the supplicant and the port is a trunk member, the port will be blocked. If the port is later removed from the trunk, the port will try to authenticate the supplicant. If authentication is successful, the port becomes unblocked. Similarly, if the supplicant is authenticated and later the port becomes a trunk member, the port will be blocked. If the port is then removed from the trunk, it tries to re-authenticate the supplicant. If successful, the port becomes unblocked.
- To help maintain security, 802.1X and LACP cannot both be enabled on the same port. If you try to configure 802.1X on a port already configured for LACP (or the reverse) you will see a message similar to the following:

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

**Note on 802.1X
and LACP**

To help maintain security, the switch does not allow 802.1X and LACP to both be enabled at the same time on the same port. Refer to “802.1X Operating Messages” on page 8-48

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

Do These Steps Before You Configure 802.1X Operation

1. Configure a local username and password on the switch for both the Operator (login) and Manager (enable) access levels. (While this may or may not be required for your 802.1X configuration, ProCurve recommends that you use a local username and password pair at least until your other security measures are in place.)
2. Determine which ports on the switch you want to operate as authenticators and/or supplicants, and disable LACP on these ports. (See the “Note on 802.1X and LACP” on page 8-11.)
3. Determine whether to use the optional 802.1X Open VLAN mode for clients that are not 802.1X-aware; that is, for clients that are not running 802.1X supplicant software. (This will require you to provide downloadable software that the client can use to enable an authentication session.) For more on this topic, refer to “802.1X Open VLAN Mode” on page 8-21.
4. For each port you want to operate as a supplicant, determine a username and password pair. You can either use the same pair for each port or use unique pairs for individual ports or subgroups of ports. (This can also be the same local username/password pair that you assign to the switch.)
5. Unless you are using only the switch’s local username and password for 802.1X authentication, configure at least one RADIUS server to authenticate access requests coming through the ports on the switch from external supplicants (including switch ports operating as 802.1X supplicants). You can use up to three RADIUS servers for authentication; one primary and two backups. Refer to the documentation provided with your RADIUS application.

Overview: Configuring 802.1X Authentication on the Switch

This section outlines the steps for configuring 802.1X on the switch. For detailed information on each step, refer to “RADIUS Authentication and Accounting” on page 5-1 or “Configuring Switch Ports To Operate As Suppliants for 802.1X Connections to Other Switches” on page 8-34.

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 8-15.
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 8-21.
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 8-19.
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 8-20.
5. Enable 802.1X authentication on the switch. See page 8-15.
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.

Configuring Port-Based Access Control (802.1X)

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

7. If you are using Port Security on the switch, configure the switch to allow only 802.1X access on ports configured for 802.1X operation, and (if desired) the action to take if an unauthorized device attempts access through an 802.1X port. See page 8-32.
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 8-34.)

Configuring Switch Ports as 802.1X Authenticators

802.1X Authentication Commands	Page
[no] aaa port-access authenticator < [ethernet] < <i>port-list</i> >	8-15
[control quiet-period tx-period supplicant-timeout server-timeout max-requests reauth-period auth-vid unauth-vid initialize reauthenticate clear-statistics]	8-15
aaa authentication port-access < local eap-radius chap-radius >	8-19
[no] aaa port-access authenticator active	8-15
[no] port-security [ethernet] < <i>port-list</i> > learn-mode port-access	8-32
802.1X Open VLAN Mode Commands	8-21
802.1X Supplicant Commands	8-34
802.1X-Related Show Commands	8-38
RADIUS server configuration	8-20

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. (Actual 802.1X operation does not commence until you perform step 5 on page 8-13 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.

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

Enables specified ports to operate as 802.1X authenticators with current per-port authenticator configuration. To activate configured 802.1X operation, you must enable 802.1X authentication. Refer to “5. Enable 802.1X Authentication on the switch” on page 8-13.

[control < authorized | auto | unauthorized >]

Controls authentication mode on the specified port:

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

auto (the default): *The device connected to the port must support 802.1X authentication and provide valid credentials in order to get network access. (You have the option of using the Open VLAN mode to provide a path for clients without 802.1X supplicant software to download this software and begin the authentication process. Refer to “802.1X Open VLAN Mode” on page 8-21.)*

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

[quiet-period < 0 - 65535 >]

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

[tx-period < 0 - 65535 >]

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

[supplicant-timeout < 1 - 300 >]

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

aaa port-access authenticator < port-list >

[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 < 1 - 9999999 >]

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

[unauth-vid < vlan-id >]

Configures an existing static VLAN to be the Unauthorized-Client VLAN. This enables you to provide a path for clients without supplicant software to download the software and begin an authentication session. Refer to “802.1X Open VLAN Mode” on page 8-21.

[auth-vid < vid >]

Configures an existing, static VLAN to be the Authorized-Client VLAN. Refer to “802.1X Open VLAN Mode” on page 8-21.

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.*

***Note:** If a specified port is configured with **control authorized** and **port-security**, and the port has learned an authorized address, the port will remove this address and learn a new one from the first packet it receives.*

[reauthenticate]

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

[clear-statistics]

Clears authenticator statistics counters.

3. Configure the 802.1X Authentication Method

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

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

Determines the type of RADIUS authentication to use.

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

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

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

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

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

Status and Counters - Authentication Information

Login Attempts : 3

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

Configuration command for EAP-RADIUS authentication.

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

Figure 8-3. 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 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.

802.1X Open VLAN Mode

802.1X Authentication Commands	page 8-15
802.1X Supplicant Commands	page 8-35
802.1X Open VLAN Mode Commands	
[no] aaa port-access authenticator [e] < port-list >	page 8-30
[auth-vid < vlan-id >]	
[unauth-vid < vlan-id >]	
802.1X-Related Show Commands	page 8-38
RADIUS server configuration	pages 8-20

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

Introduction

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

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

The 802.1X Open VLAN mode solves this problem by temporarily suspending the port’s static, 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. Following client authentication, the port drops its temporary (untagged) membership in the Unauthorized-Client VLAN and joins (or rejoins) *one* of the following as an *untagged* member:

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.

If the port is not configured for any of the above, then it must be a tagged member of at least one static VLAN. If the client is capable of operating with that tagged VLAN, then it receives access to the VLAN. Otherwise, the connection fails.

Note

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 unauthenticated, friendly clients will need access to some services before being authenticated.
- **Authorized-Client VLAN:** Configure this VLAN for authenticated clients when the port is not statically configured as an untagged member of a VLAN you want clients to use, or when the port is statically configured as an untagged member of a VLAN you do not want clients to use. (A port can be configured as untagged on only one VLAN. When an Authorized-Client VLAN is configured, it will always be untagged and will block the port from using a statically configured, untagged membership in another VLAN.) 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 8-1. 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	<ul style="list-style-type: none"> • When the port detects a client, it automatically becomes an untagged member of this VLAN. If you previously configured the port as a static, tagged member of the VLAN, membership temporarily changes to untagged while the client remains unauthenticated. • If the port already has a statically configured, untagged membership in another VLAN, then the port temporarily closes access to this other VLAN while in the Unauthorized-Client VLAN. • To limit security risks, the network services and access available on the Unauthorized-Client VLAN should include only what a client needs to enable an authentication session. If the port is statically configured as a tagged member of any other VLANs, access to these VLANs is blocked while the port is a member of the Unauthorized-Client VLAN.
Authorized-Client VLAN	<ul style="list-style-type: none"> • After the client is authenticated, the port drops membership in the Unauthorized-Client VLAN and becomes an untagged member of this VLAN. 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.

802.1X Per-Port Configuration	Port Response
-------------------------------	---------------

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

- When the port detects a client, it automatically becomes an untagged member of this VLAN. To limit security risks, the network services and access available on this VLAN should include only what a client needs to enable an authentication session. If the port is statically configured as an untagged member of another VLAN, the switch temporarily removes the port from membership in this other VLAN while membership in the Unauthorized-Client VLAN exists.
- After the client is authenticated, and if the port is statically configured as an untagged member of another VLAN, the port's access to this other VLAN is restored.

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:

- 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 <i>Authorized-Client</i> or <i>Unauthorized-Client</i> VLANs	These must be configured on the switch before you configure an 802.1X authenticator port to use them. (Use the vlan < vlan-id > command or the VLAN Menu screen in the Menu interface.)
VLAN Assignment Received from a RADIUS Server	If the RADIUS server specifies a VLAN for an authenticated supplicant connected to an 802.1X authenticator port, this VLAN assignment overrides any Authorized-Client VLAN assignment configured on the authenticator port. This is because both VLANs are untagged, and the switch allows only one untagged VLAN membership per-port. For example, suppose you configured port A4 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 8-22.)
Temporary VLAN Membership During a Client Session	<ul style="list-style-type: none"> • Port membership in a VLAN assigned to operate as the Unauthorized-Client VLAN is temporary, and ends when the client receives authentication or the client disconnects from the port, whichever is first. • Port membership in a VLAN assigned to operate as the Authorized-Client VLAN is also temporary, and ends when the client disconnects from the port. If a VLAN assignment from a RADIUS server is used instead, the same rule applies.
Effect of Unauthorized-Client VLAN session on untagged port VLAN membership	<ul style="list-style-type: none"> • When an unauthenticated client connects to a port that is already configured with a static, untagged VLAN, the switch temporarily moves the port to the Unauthorized-Client VLAN (also untagged). (While the Unauthorized-Client VLAN is in use, the port does not access the static, untagged VLAN.) • When the client either becomes authenticated or disconnects, the port leaves the Unauthorized-Client VLAN and reacquires its untagged membership in the statically configured VLAN.

Configuring Port-Based Access Control (802.1X)
802.1X Open VLAN Mode

Condition	Rule
Effect of Authorized-Client VLAN session on untagged port VLAN membership.	<ul style="list-style-type: none">• When a client becomes authenticated on a port that is already configured with a static, untagged VLAN, the switch temporarily moves the port to the Authorized-Client VLAN (also untagged). While the Authorized-Client VLAN is in use, the port does not have access to the statically configured, untagged VLAN.• When the authenticated client disconnects, the switch removes the port from the Authorized-Client VLAN and moves it back to the untagged membership in the statically configured VLAN. (After client authentication, the port resumes any tagged VLAN memberships for which it is already configured. For details, refer to the Note on page 8-22.)
Multiple Authenticator Ports Using the Same Unauthorized-Client and Authorized-Client VLANs	<p>You can use the same static VLAN as the Unauthorized-Client VLAN for all 802.1X authenticator ports configured on the switch. Similarly, you can use the same static VLAN as the Authorized-Client VLAN for all 802.1X authenticator ports configured on the switch.</p> <p>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.</p>
Effect of Failed Client Authentication Attempt	<p>When there is an Unauthorized-Client VLAN configured on an 802.1X authenticator port, an unauthorized client connected to the port has access only to the network resources belonging to the Unauthorized-Client VLAN. This access continues until the client disconnects from the port. (If there is no Unauthorized-Client VLAN configured on the authenticator port, the port simply blocks access for any unauthorized client that cannot be authenticated.)</p>
IP Addressing for a Client Connected to a Port Configured for 802.x Open VLAN Mode	<p>A client can either acquire an IP address from a DHCP server or have a preconfigured, manual IP address before connecting to the switch.</p>
802.1X Supplicant Software for a Client Connected to a Port Configured for 802.1X Open VLAN Mode	<p>A friendly client, without 802.1X supplicant software, connecting to an authenticator port must be able to download this software from the Unauthorized-Client VLAN before authentication can begin.</p>

Note

If you use the same VLAN as the Unauthorized-Client VLAN for all authenticator ports, unauthenticated clients on different ports can communicate with each other. However, in this case, you can improve security between authenticator ports by using the switch's Source-Port filter feature. For example, if you are using ports B1 and B2 as authenticator ports on the same Unauthorized-Client VLAN, you can configure a Source-Port filter on B1 to drop all packets from B2 and the reverse.

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 8-1 on page 8-23 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 A5 is an untagged member of VLAN 1 (the default VLAN).
- ii. You configure port A5 as an 802.1X authenticator port.
- iii. You configure port A5 to use an Authorized-Client VLAN.

Then, if a client connects to port A5 and is authenticated, port A5 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 8-15.) This setting requires a client to support 802.1X authentication (with 802.1X supplicant operation) and to provide valid credentials to get network access.

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

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

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

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

Determines the type of RADIUS authentication to use.

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

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

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

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

Syntax: radius host < ip-address >

Adds a server to the RADIUS configuration.

[key < server-specific key-string >]

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

radius-server key < global key-string >

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

4. Activate authentication on the switch.

Syntax: aaa port-access authenticator active

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

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

Note

If you want to implement the optional port security feature on the switch, you should first ensure that the ports you have configured as 802.1X authenticators operate as expected. Then refer to “Option For Authenticator Ports: Configure Port-Security To Allow Only 802.1X Devices” on page 8-32.

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 8-27.

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 A10 - A20 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

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

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

```
ProCurve(config)# aaa port-access authenticator a10-a20
```

Configures ports A10 - A20 as 802.1 authenticator ports.

```
ProCurve(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.

```
ProCurve(config)# aaa port-access authenticator e a10-a20 unauth-vid 80
```

Configures ports A10 - A20 to use VLAN 80 as the Unauthorized-Client VLAN.

```
ProCurve(config)# aaa port-access authenticator e a10-a20 auth-vid 81
```

Configures ports A10 - A20 to use VLAN 81 as the Authorized-Client VLAN.

```
ProCurve(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 8-40.

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.

Configuring Port-Based Access Control (802.1X)

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

- If an authenticated client loses authentication during a session in 802.1X Open VLAN mode, the port VLAN membership reverts back to the Unauthorized-Client VLAN. If there is no Unauthorized-Client VLAN configured, then the client loses access to the port until it can reauthenticate itself.

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 A10 for 802.1X authenticator operation and display the result:

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

**Note on
Blocking a Non-
802.1X Device**

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

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

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

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

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

Enables 802.1X authentication on the port.

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

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

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

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

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

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

802.1X Authentication Commands page 8-15

802.1X Supplicant Commands

[no] aaa port-access < supplicant > [ethernet] < port-list > page 8-35

[auth-timeout | held-period | start-period | max-start | initialize |
identity | secret | clear-statistics] page 8-36

802.1X-Related Show Commands page 8-38

RADIUS server configuration pages 8-20

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

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

- Switch “A” has port A1 configured for 802.1X supplicant operation
- You want to connect port A1 on switch “A” to port B5 on switch “B”.




Figure 8-4. Example of Supplicant Operation

1. When port A1 on switch “A” is first connected to a port on switch “B”, or if the ports are already connected and either switch reboots, port A1 begins sending start packets to port B5 on switch “B”.
 - If, after the supplicant port sends the configured number of start request packets, it does not receive a response, it assumes that switch “B” is not 802.1X-aware, and transitions to the authenticated state. If switch “B” is operating properly and is not 802.1X-aware, then the link should begin functioning normally, but without 802.1X security.
 - If, after sending one or more start request packets, port A1 receives a request packet from port B5, then switch “B” is operating as an 802.1X authenticator. The supplicant port then sends a response/ID packet. If switch “B” is configured for RADIUS authentication, it forwards this request to a RADIUS server. If switch “B” is configured for Local 802.1X authentication (page 8-19), 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 A1 on switch “A”.
3. Port A1 replies with a hash response based on its unique credentials. Switch “B” forwards this response to the RADIUS server.
4. The RADIUS server then analyzes the response and sends either a “success” or “failure” packet back through switch “B” to port A1.
 - A “success” response unblocks port B5 to normal traffic from port A1.
 - A “failure” response continues the block on port B5 and causes port A1 to wait for the “held-time” period before trying again to achieve authentication through port B5.

Note

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

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

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

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

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

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

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

[identity < username >]

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

[secret]

Enter secret: < password >

Repeat secret: < password >

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

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

[auth-timeout < 1 - 300 >]

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

[max-start < 1 - 10 >]

Defines the maximum number of times the supplicant port requests authentication. See step 1 on page 8-35 for a description of how the port reacts to the authenticator response. (Default: 3).

[held-period < 0 - 65535 >]

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

[start-period < 1 - 300 >]

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

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

[initialize]

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

[clear-statistics]

Clears and restarts the 802.1X supplicant statistics counters.

Displaying 802.1X Configuration, Statistics, and Counters

802.1X Authentication Commands	page 8-15
802.1X Supplicant Commands	page 8-34
802.1X Open VLAN Mode Commands	page 8-21
802.1X-Related Show Commands	
show port-access authenticator	below
show port-access supplicant	page 8-43
Details of 802.1X Mode Status Listings	page 8-40
RADIUS server configuration	pages 8-20

Show Commands for Port-Access Authenticator

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

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

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

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

config [[e] < port-list >]

Shows:

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

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

statistics [[e] < port-list >]

Shows:

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

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

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

Shows:

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

*Also, for each port, the "User" column lists the user name the supplicant included in its response packet. (For the switch, this is the **identity** setting included in the **supplicant** command—page 8-36.) Does not display data for a specified port that is not enabled as an authenticator.*

Viewing 802.1X Open VLAN Mode Status

You can examine the switch's current VLAN status by using the **show port-access authenticator** and **show vlan < vlan-id >** commands as illustrated in this section. Figure 8-5 shows an example of **show port-access authenticator** output, and table 8-1 describes the data that this command displays. Figure 8-6 shows related VLAN data that can help you to see how the switch is using statically configured VLANs to support 802.1X operation.

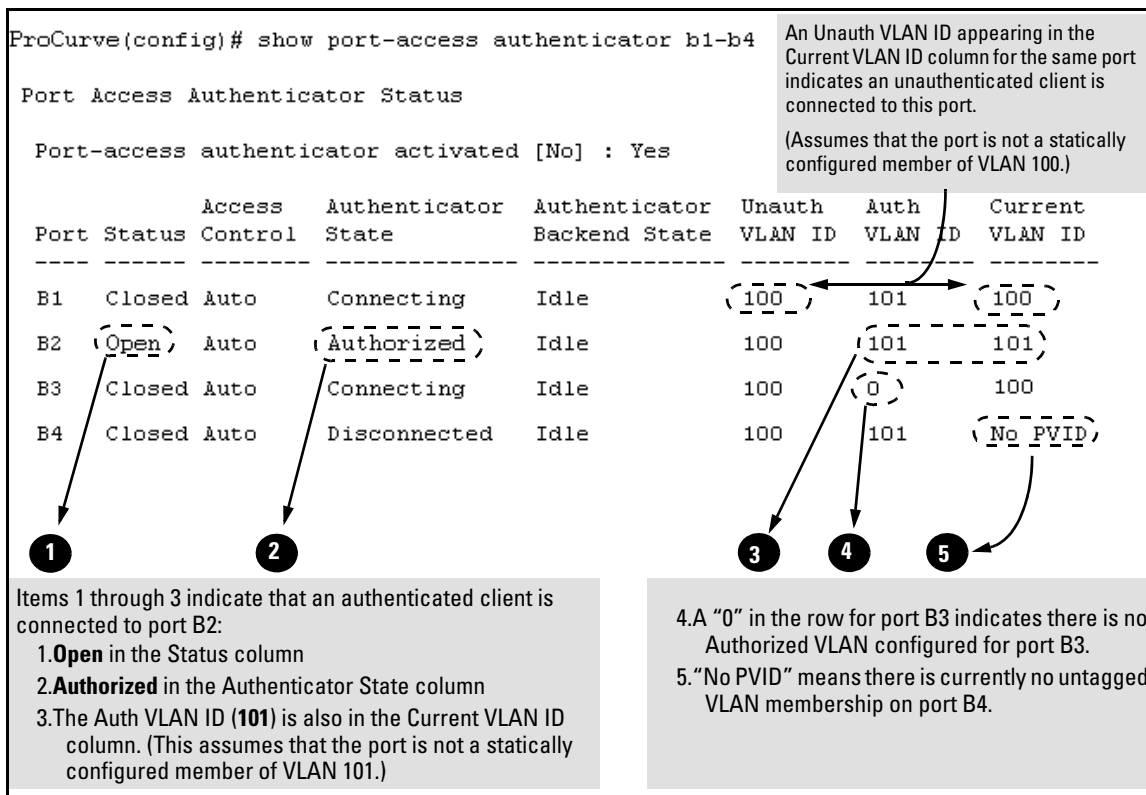


Figure 8-5. Example Showing Ports Configured for Open VLAN Mode

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 A12 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 8-2. Open VLAN Mode Status

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

Configuring Port-Based Access Control (802.1X)
 Displaying 802.1X Configuration, Statistics, and Counters

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.

```

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

Port Information Mode      Unknown VLAN Status
-----
A1             Untagged Learn      Up
A2             Untagged Learn      Up
A3             Untagged Learn      Up
A4             Untagged Learn      Up
B2             Untagged Learn      Up
B4             Tagged Learn       Up
B5             Untagged Learn      Down
.             .
.             .
.             .
B23            Untagged Learn      Up
B24            Untagged Learn      Up

Overridden Port VLAN configuration
Port Mode
-----
B1  Untagged
B3  Untagged
    
```

Note that ports B1 and B3 are not in the upper listing, but are included under "Overridden Port VLAN configuration". This shows that static, untagged VLAN memberships on ports B1 and B3 have been overridden by temporary assignment to the authorized or unauthorized VLAN. Using the **show port-access authenticator < port-list >** command shown in figure 8-5 provides details.

Figure 8-6. Example of Showing a VLAN with Ports Configured for Open VLAN Mode

Show Commands for Port-Access Supplicant

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

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

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

Connecting - Starting authentication.

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

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

Authenticating - Authentication is in progress.

Held - Authenticator sent notice of failure. The supplicant port is waiting for the authenticator's held-period (page 8-36).

For descriptions of the supplicant parameters, refer to “Configuring a Supplicant Switch Port” on page 8-36.

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

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

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

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

Thus, if the supplicant's link to the authenticator fails, the supplicant retains the transaction statistics it most recently received until one of the above events occurs. Also, if you move a link with an authenticator from one

supplicant port to another without clearing the statistics data from the first port, the authenticator's MAC address will appear in the supplicant statistics for both ports.

How RADIUS/802.1X Authentication Affects VLAN Operation

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

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

For example, suppose that a RADIUS-authenticated, 802.1X-aware client on port A2 requires access to VLAN 22, but VLAN 22 is configured for no access on port A2, and VLAN 33 is configured as untagged on port A2:

```

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

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

Actions->  Cancel  Edit    Save    Help
Cancel changes and return to previous screen.
Use arrow keys to change action selection and <Enter> to execute.
    
```

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

Figure 8-7. Example of an Active VLAN Configuration

In figure 8-7, 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 A2 for the duration of the session.
- VLAN 33 becomes unavailable to port A2 for the duration of the session (because there can be only one untagged VLAN on any port).

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

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

Configuring Port-Based Access Control (802.1X)
 How RADIUS/802.1X Authentication Affects VLAN Operation

```

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

Port Information Mode      Unknown VLAN Status
-----
A1                Tagged   Learn           Up
A2                (802.1X)  Learn           Up
A4                Tagged   Learn           Up
.                 .         .               .
.                 .         .               .
.                 .         .               .

Overridden Port VLAN configuration

Port  Mode
----  ---
A2    (No)
  
```

This entry shows that port A2 is temporarily untagged on VLAN 22 for an 802.1X session. This is to accommodate an 802.1X client's access, authenticated by a RADIUS server, where the server included an instruction to put the client's access on VLAN 22.

Note: With the current VLAN configuration (figure 8-7), the only time port A2 appears in this **show vlan 22** listing is during an 802.1X session with an attached client. Otherwise, port A2 is not listed.

Figure 8-8. The Active Configuration for VLAN 22 Temporarily Changes for the 802.1X Session

- With the preceding in mind, since (static) VLAN 33 is configured as untagged on port A2 (see figure 8-7), and since a port can be untagged on only one VLAN, port A2 loses access to VLAN 33 for the duration of the 802.1X session involving VLAN 22. You can verify the temporary loss of access to VLAN 33 with the **show vlan 33** command.

```

ProCurve# show vlan 33
Status and Counters - VLAN Information - Ports - VLAN 33
802.1Q VLAN ID : 33
Name          : VLAN_33
Status       : Static

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

Overridden Port VLAN configuration

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

Even though port A2 is configured as Untagged on (static) VLAN 33 (see figure 8-7), it does not appear in the VLAN 33 listing while the 802.1X session is using VLAN 22 in the Untagged status. However, after the 802.1X session with VLAN 22 ends, the active configuration returns port A2 to VLAN 33.

Figure 8-9. The Active Configuration for VLAN 33 Temporarily Drops Port 22 for the 802.1X Session

When the 802.1X client's session on port A2 ends, the port discards the temporary untagged VLAN membership. At this time the static VLAN actually configured as untagged on the port again becomes available. Thus, when the RADIUS-authenticated 802.1X session on port A2 ends, VLAN 22 access on port A2 also ends, and the untagged VLAN 33 access on port A2 is restored.

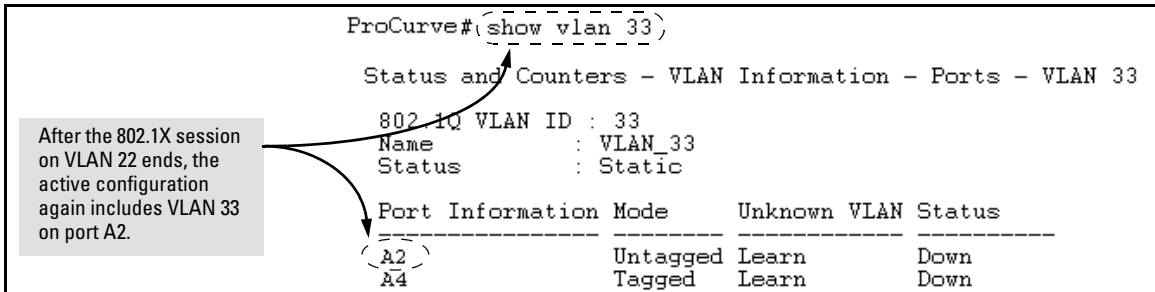


Figure 8-10. The Active Configuration for VLAN 33 Restores Port A2 After the 802.1X Session Ends

Notes

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

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

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

Messages Related to 802.1X Operation

Table 8-1. 802.1X Operating Messages

Message	Meaning
Port < port-list > is not an authenticator.	The ports in the port list have not been enabled as 802.1X authenticators. Use this command to enable the ports as authenticators: <pre>ProCurve(config)# aaa port-access authenticator e 10</pre>
Port < port-list > is not a supplicant.	Occurs when there is an attempt to change the supplicant configuration on a port that is not currently enabled as a supplicant. Enable the port as a supplicant and then make the desired supplicant configuration changes. Refer to "Enabling a Switch Port To Operate as a Supplicant" on page 8-35.
No server(s) responding.	This message can appear if you configured the switch for EAP-RADIUS or CHAP-RADIUS authentication, but the switch does not receive a response from a RADIUS server. Ensure that the switch is configured to access at least one RADIUS server. (Use show radius .) If you also see the message <code>Can't reach RADIUS server < x.x.x.x ></code> , try the suggestions listed for that message (page 5-32).
LACP has been disabled on 802.1X port(s).	To maintain security, LACP is not allowed on ports configured for 802.1X authenticator operation. If you configure port security on a port on which LACP (active or passive) is configured, the switch removes the LACP configuration, displays a notice that LACP is disabled on the port(s), and enables 802.1X on that port.
Error configuring port < port-number >: LACP and 802.1X cannot be run together.	Also, the switch will not allow you to configure LACP on a port on which port access (802.1X) is enabled.

Configuring and Monitoring Port Security

Contents

Overview	9-2
Basic Operation	9-2
Blocking Unauthorized Traffic	9-3
Trunk Group Exclusion	9-4
Planning Port Security	9-5
Port Security Command Options and Operation	9-6
Retention of Static MAC Addresses	9-10
Displaying Current Port Security Settings	9-10
Configuring Port Security	9-12
MAC Lockdown	9-17
Differences Between MAC Lockdown and Port Security	9-19
Deploying MAC Lockdown	9-21
MAC Lockout	9-25
Port Security and MAC Lockout	9-27
Web: Displaying and Configuring Port Security Features	9-27
Reading Intrusion Alerts and Resetting Alert Flags	9-28
Notice of Security Violations	9-28
How the Intrusion Log Operates	9-29
Keeping the Intrusion Log Current by Resetting Alert Flags	9-29
Using the Event Log To Find Intrusion Alerts	9-35
Web: Checking for Intrusions, Listing Intrusion Alerts, and Resetting Alert Flags	9-36
Operating Notes for Port Security	9-36

Overview

Feature	Default	Menu	CLI	Web
Displaying Current Port Security	n/a	—	page 9-10	page 9-27
Configuring Port Security	disabled	—	page 9-12	page 9-27
Intrusion Alerts and Alert Flags	n/a	page 9-35	page 9-33	page 9-36

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

Note

This feature does not prevent intruders from receiving broadcast and multi-cast traffic.

Basic Operation

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

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

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

- Alert flags that are captured by network management tools
- Alert Log entries in the switch's web browser interface
- Event Log entries in the console interface
- Intrusion Log entries in either the menu interface, CLI, or web browser interface

For any port, you can configure the following:

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

Blocking Unauthorized Traffic

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

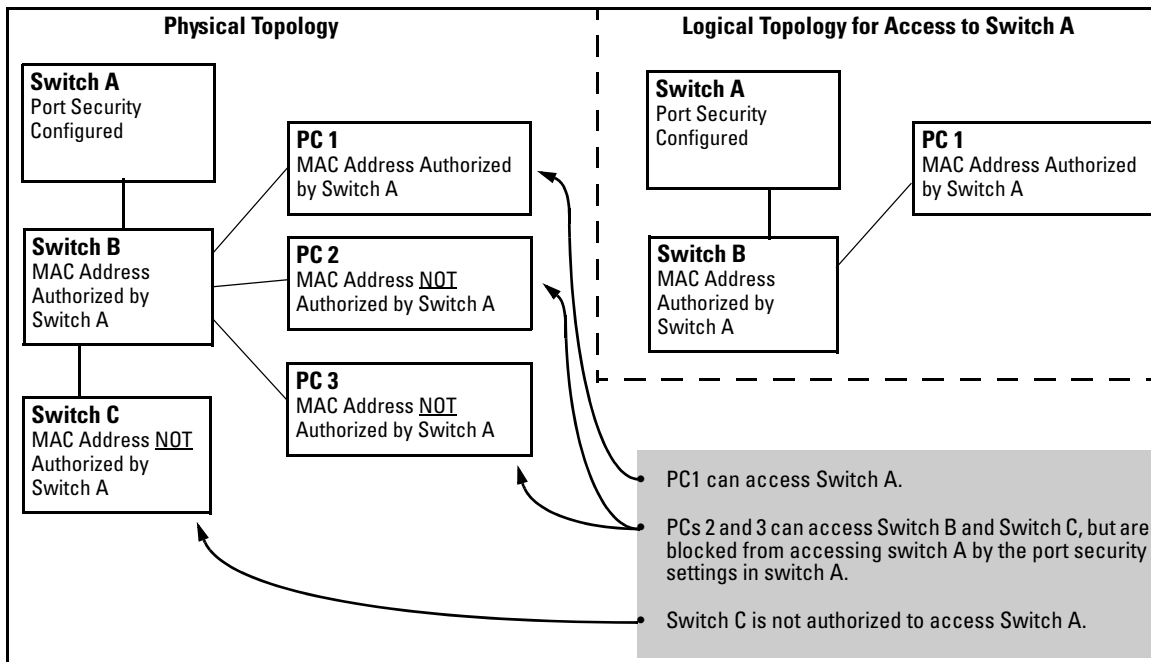


Figure 9-1. Example of How Port Security Controls Access

Note

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

Trunk Group Exclusion

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

Planning Port Security

1. Plan your port security configuration and monitoring according to the following:
 - a. On which ports do you want port security?
 - b. Which devices (MAC addresses) are authorized on each port and how many devices do you want to allow per port (up to 8)?
 - c. Within the devices-per-port limit, do you want to let the switch automatically accept devices it detects on a port, or do you want it to accept only the devices you explicitly specify? (For example, if you allow three devices on a given port, but specify only one MAC address for that port, do you want the switch to automatically accept the first two additional devices it detects, or not?)
 - d. 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.
 - e. 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	9-11
port-security	9-12
< [ethernet] <i>port-list</i> >	9-12
[learn-mode]	9-12
[address-limit]	9-12
[mac-address]	9-12
[action]	9-12
[clear-intrusion-flag]	9-12
no port-security	9-12

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

Note

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

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

learn-mode < continuous | static | configured | port-access >

Continuous (Default): *Appears in the factory-default setting or when you execute **no port-security**. Allows the port to learn addresses from inbound traffic from any device(s) to which it is connected. In this state, the port accepts traffic from any device(s) to which it is connected. Addresses learned this way appear in the switch and port address tables and age out according to the **MAC Age Interval** in the System Information configuration screen of the Menu interface or the **show system-information** listing.*

Static: *The static-learn option enables you to use the **mac-address** parameter to specify the MAC addresses of the devices authorized for a port, and the **address-limit** parameter to specify the number of MAC addresses authorized for the port. You can authorize specific devices for the port, while still allowing the port to accept other, non-specified devices until the port reaches the configured address limit. That is, if you enter fewer MAC addresses than you authorized, the port fills the remainder of the address allowance with MAC addresses it automatically learns. For example, if you specify three authorized devices, but enter 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 authorize MAC address **0060b0-880a80** on port A4.
- You allow three devices on port A4, but the port detects these MAC addresses:
 1. **080090-1362f2**
 2. **00f031-423fc1**
 3. **080071-0c45a1**
 4. **0060b0-880a80** (the authorized address.)

Port A4 then has the following list of authorized addresses:

- 080090-1362f2** (The first address detected.)
- 00f031-423fc1** (The second address detected.)
- 0060b0-880a80** (The authorized address.)

*The remaining MAC address, **080071-0c45a1**, is an intruder. See also “Retention of Static Addresses” on page 9-10.*

Caution: When you use **learn-mode static** with a device limit greater than the number of MAC addresses you specify with **mac-address**, an unwanted device can become “authorized”. This can occur because the port, in order to fulfill the number of devices allowed by **address-limit**, automatically adds devices it detects until it reaches the specified limit.

Syntax: port-security [e] < port-list > (- **Continued** -)

learn-mode < continuous | static | configured | port-access >

Configured: *The static-configured option operates the same as the static-learn option on the preceding page, except that it does not allow the switch to accept non-specified addresses to reach the address limit. Thus, if you configure an address limit of 3, but only configure two MAC addresses, the switch will handle as intruders all non-specified MAC addresses it detects.*

Port-Access: *Enables you to use Port Security with (802.1X) Port-Based Access Control. Refer to “Configuring Port-Based Access Control (802.1X)” on page 8-1.*

address-limit < integer >

*When Learn Mode is set to **static** (static-learn) or **configured** (static-configured), this parameter specifies the number of authorized devices (MAC addresses) to allow. Default: 1; Range: 1 to 8.*

mac-address < mac-addr >

*Available for **static** (static-learn and configured-learn) modes. Allows up to eight authorized devices (MAC addresses) per port, depending on the value specified in the **address-limit** parameter.*

- *If you use **mac-address** with **learn-mode configured**, but enter fewer devices than you specified in the **address-limit** field, the port accepts only the devices you specified with **mac-address**. (See the **Note**, above.)*
- *If you use **mac-address** with **learn-mode static**, but enter fewer devices than you specified in the **address-limit** field, the port accepts the specified devices AND as many other devices as it takes to reach the device limit.*

Syntax: port-security [e] <port-list> (- Continued -)

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

Specifies whether an SNMP trap is sent to a network management station. Operates when:

- Learn mode is set to **learn-mode static** (*static-learn*) or **learn-mode configured** (*static-configured*) and the port detects an unauthorized device.
- Learn mode is set to **learn-mode continuous** and there is a MAC address change on a port.

none (the default): Prevents an SNMP trap from being sent.

send alarm: Causes the switch to send an SNMP trap to a network management station.

send-disable: Available only with **learn-mode configured** and **learn-mode static**. 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 will block further intruders, but the switch will not disable the port again until you reset the intrusion flag. See the **Note** on page 9-30.

For information on configuring the switch for SNMP management, refer to the Management and Configuration Guide for your switch.

clear-intrusion-flag

Clears the intrusion flag for a specific port. (Refer to "Reading Intrusion Alerts and Resetting Alert Flags" on page 9-28.)

Retention of Static MAC Addresses

Learned MAC Addresses

In the following two cases, a port in Static learn mode (**learn-mode static**) 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 with **learn-mode static** in both the startup-config file and the running-config file (by executing **write memory**).
- The port learns a MAC address after you configure the port with **learn-mode static** 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.

Assigned/Authorized MAC Addresses

If you manually assign a MAC address (using **mac-address < mac-addr >**) and then execute **write memory**, the assigned MAC address remains in memory unless removed by one of the methods described below.

Removing Learned and Assigned Static MAC Addresses

To remove a static MAC address, 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.

Displaying Current Port Security Settings

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

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

Using the CLI To Display Port Security Settings.

Syntax: show port-security
show port-security [e] <port number>
show port-security [e] [<port number>-<port number>] . . [<port number>]

Without port parameters, **show port-security** displays operating control settings for all ports on a switch. For example:

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

Figure 9-2. Example Port Security Listing (Ports A7 and A8 Show the Default Setting)

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

```
ProCurve(config)# show port-security A3
Port Security
  Port : A3
  Learn Mode [Continuous]: Static      Address Limit[1]:
  Action [None]: Send Alarm
  Authorized Addresses
  -----
  00906d-fdcc00
```

Figure 9-3. Example of the Port Security Configuration Display for a Single Port

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

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

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 <continuous | static | configured | port-access >]
[address-limit <integer >]
[mac-address <mac-addr >] [<mac-addr >... <mac-addr >]
[action <none | send-alarm | send-disable >]
[clear-intrusion-flag]

(For the **configured** option, above, refer to the **Note** on page 9-6.

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

Specifying Authorized Devices and Intrusion Responses

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

```
ProCurve(config)# port-security a1 learn-mode static  
action send-disable
```

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

```
ProCurve(config)# port-security a1 learn-mode static  
mac-address 0c0090-123456 action send-disable
```

This example configures port A5 to:

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

```
ProCurve(config)# port-security a5 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 reset the switch to its factory-default configuration. You can “turn off” device authorization on a port by configuring the port to **continuous** Learn Mode, but subsequently reconfiguring the port to **static** Learn Mode restores the configured device authorization.

Learn-Mode Configured. This option allows only MAC addresses specifically configured with **learn-mode configured mac-address < mac-address >**, and does not automatically learn non-specified MAC addresses learned from the network. This example configures port A1 to:

- Allow only a MAC address of 0c0090-123456 as the authorized device
- Reserve the option for adding two more specified MAC addresses at a later time without having to change the address-limit setting.
- Send an alarm to a management station if an intruder is detected on the port.

```
ProCurve(config)# port-security A1 learn-mode configured  
mac-address 0c0090-123456 address-limit 3 action send-  
disable
```

Adding a MAC Address to an Existing Port List

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 either **static** or **configured** and the Authorized Addresses list is not already full* (as deter-

mined by the current **address-limit** value). For example, suppose port A1 allows two authorized devices, but has only one device in its Authorized Address list:

```
ProCurve(config)# show port-security a1
Port Security
Port : A1
Learn Mode [Continuous] : Static      Address Limit [1] : 2
Action [None] : None

Authorized Addresses
-----
0c0090-123456
```

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

The Address Limit has not been reached.

Figure 9-4. Example of Adding an Authorized Device to a Port

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

```
ProCurve(config)# port-security a1 mac-address 0c0090-456456
```

After executing the above command, the security configuration for port A1 appears as:

```
ProCurve(config)# show port-security a1
Port Security
Port : A1
Learn Mode [Continuous] : Static      Address Limit [1] : 2
Action [None] : None

Authorized Addresses
-----
0c0090-123456
0c0090-456456
```

The Address Limit has been reached.

Figure 9-5. Example of Adding a Second Authorized Device to a Port

Note

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. If you change a port from static to continuous learn mode, the port retains in memory any authorized addresses it had while in static mode. If you subsequently

attempt to convert the port back to static mode with the same authorized address(es), the `Inconsistent value` message appears because the port already has the address(es) in its “Authorized” list.

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

```
ProCurve(config)# show port-security a1
Port Security
  Port : A1
  Learn Mode [Continuous] : Static   Address Limit [1]:1
  Action [None] : None

  Authorized Addresses
  -----
  0c0090-123456
```

Figure 9-6. Example of Port Security on Port A1 with an Address Limit of “1”

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

```
ProCurve(config)# port-security a1 mac-address 0c0090-456456 address-limit 2
```

Removing a Device From the “Authorized” List for a Port Configured for Learn-Mode Static. This command option removes unwanted devices (MAC addresses) from the Authorized Addresses list. (An Authorized Address list is available for each port for which Learn Mode is currently set to “Static”. See the “MAC Address” entry in the table on 9-8.)

Caution

The **address-limit** setting controls how many MAC addresses are allowed in the Authorized Addresses list for a given port. If you remove a MAC address without also reducing the address limit by 1, the port may later detect and accept the same or another MAC address that you do not want in the Authorized Address list. Thus, if you use the CLI to remove a MAC address that is no longer authorized, you should first reduce the Address Limit (**address-limit**) integer by 1, as shown in the next example. This prevents the possibility of the same device or another device on the network from automatically being

accepted as “authorized” for that port. (You can prevent the port from learning unauthorized MAC addresses by using the **learn-mode configured** option instead of the **learn-mode static** option. Refer to the **Note** on page 9-6.)

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

Note

When you have configured the switch for **learn-mode static** operation, 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. (If you use learn-mode configured instead, the switch cannot automatically add detected devices not included in the **mac-address** configuration. Refer to the **Note** on page 9-6.)

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

```
ProCurve(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
```

Figure 9-7. Example of Two Authorized Addresses on Port A1

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

```
ProCurve(config)# port-security a1 address-limit 1
ProCurve(config)# no port-security a1 mac-address 0c0090-123456
```

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


```
ProCurve(config)# show port-security 1

Port Security

Port : 1
Learn Mode : Static      Address Limit : 1
Action: None

Authorized Addresses
-----
0c0090-456456
```

Figure 9-8. Example of Port A1 After Removing One MAC Address

MAC Lockdown

MAC Lockdown, also known as “static addressing,” is the permanent assignment 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.

Note

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: module A: Move 0001e6-1f96c0  
to A15 denied
```

```
W 10/30/03 21:33:48 maclock: module A: Move 0001e6-1f96c0  
to A15 denied
```

```
W 10/30/03 21:33:48 maclock: module A: 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 also, but message throttling is imposed on the logging on a per-module basis. 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 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.

Let's examine a good use of MAC Lockdown within a network to ensure security first.

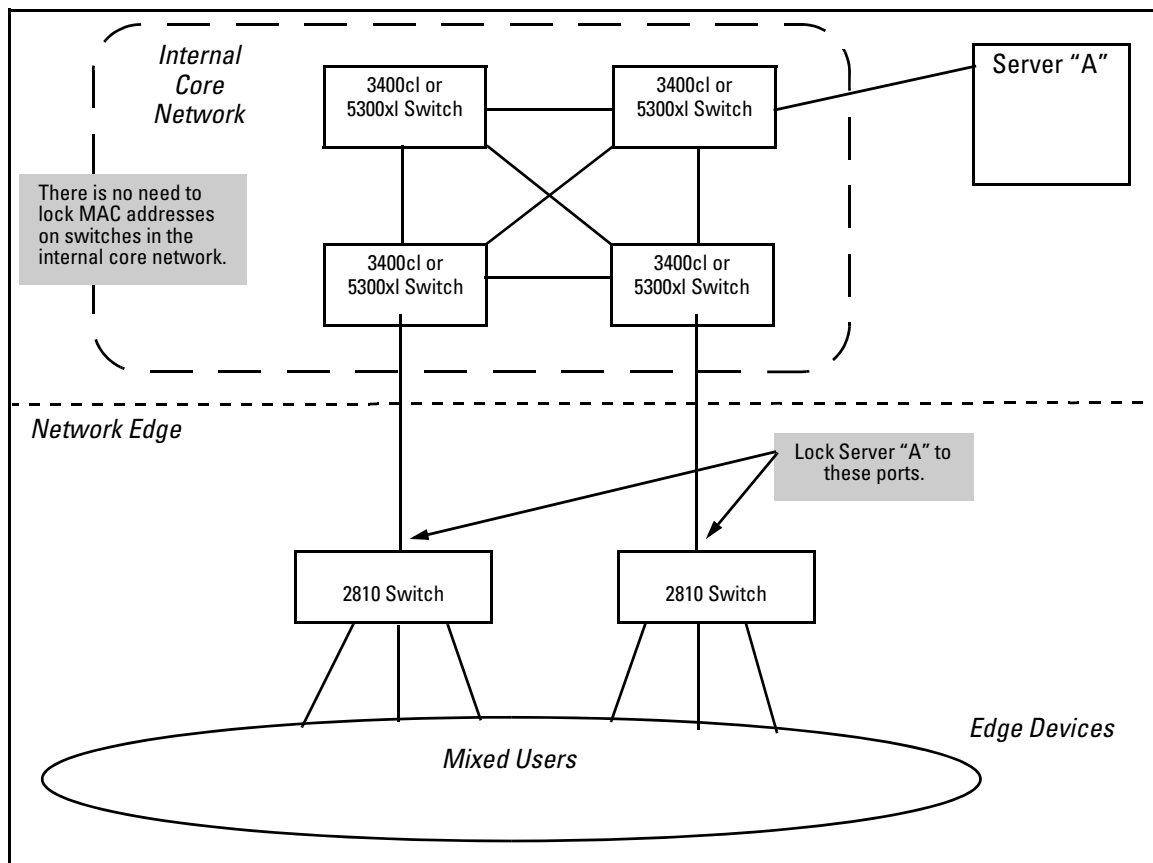


Figure 9-9. MAC Lockdown Deployed At the Network Edge Provides Security

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.

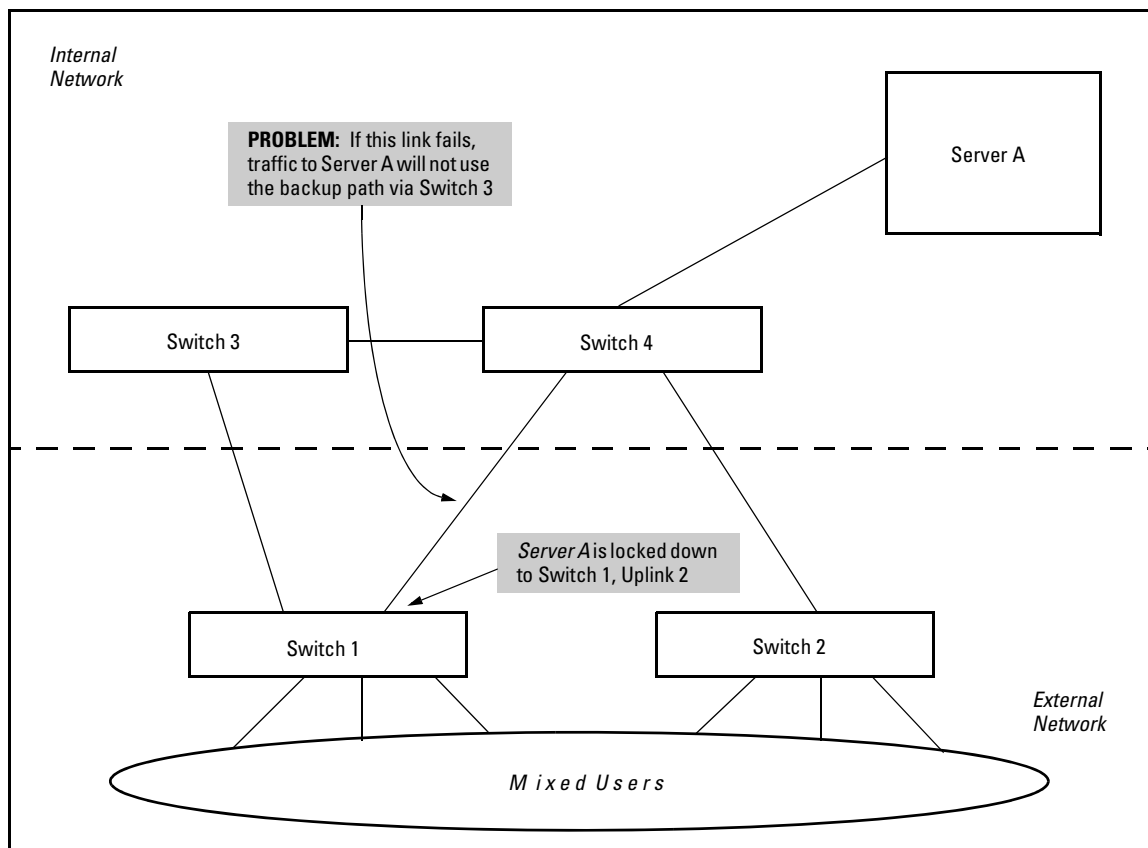


Figure 9-10. Connectivity Problems Using MAC Lockdown with Multiple Paths

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 STP or having an alternate path.

Technologies such as STP are primarily intended for an internal campus network environment in which all users are trusted. STP does not work well with MAC Lockdown.

If you deploy MAC Lockdown as shown in the Model Topology in figure 9-9 (page 9-22), you should have no problems with either security or connectivity.

Displaying status. Locked down ports are listed in the output of the **show running-config** command in the CLI. The **show static-mac** command also lists the locked down MAC addresses, as shown below.

```
ProCurve(config)# show static-mac
VLAN  MAC Address Port
  1 001083-34f8fa 9
Number of locked down MAC addresses = 1
ProCurve(config)#
```

Figure 9-11. Listing Locked Down Ports

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.

The number of MAC lockouts allowed per VLAN depends on the number of VLANs you have configured, as shown below.

Table 9-1. Number of MAC Lockouts with VLANs

Max-VLANs	Lockouts
8	50
16	25
>16	2

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 >

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)

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: module A: 0001e6-1f96c0
detected on port A15
W 10/30/03 21:35:18 maclock: module A: 0001e6-1f96c0
detected on port A15
W 10/30/03 21:35:18 maclock: module A: 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 overlogged with error messages. (Refer to "Limiting the Frequency of Log Messages" on page 9-20.)

Displaying status. Locked out ports are listed in the output of the **show running-config** command in the CLI. The **show lockout-mac** command also lists the locked out MAC addresses, as shown below.

```
ProCurve(config)# show lockout-mac
Locked Out Addresses
007347-a8fd30
Number of locked out MAC addresses = 1
ProCurve(config)#
```

Figure 9-12. Listing Locked Out Ports

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.

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 an active network management environment 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 A1 and the Intrusion Log shows two or more entries for port 1, only the most recent entry has not been acknowledged (by resetting the alert flag). The other entries give you a history of past intrusions detected on port A1.

```
ProCurve# show port-security intrusion-log
Status and Counters - Intrusion Log
Port  MAC Address          Date / Time
-----
A1    080009-e93d4f            07/03/02 21:09:34
A1    080009-21ae84            07/03/02 17:26:27
A1    080009-e93d4f prior to 07/03/02 17:18:43
```

Figure 9-13. 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 then an intruder is detected on the port, 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 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**

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

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

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

Figure 9-14. Example of Port Status Screen with Intrusion Alert on Port A3

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

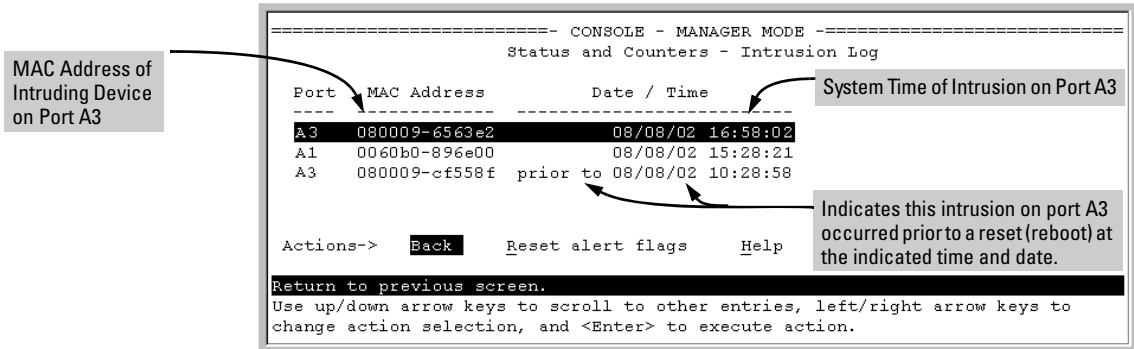


Figure 9-15. Example of the Intrusion Log Display

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

- Because the Port Status screen (figure 9-14 on page 9-30) does not indicate an intrusion for port A1, the alert flag for the intrusion on port A1 has already been reset.
- Since the switch can show only one uncleared intrusion per port, the older intrusion for port A3 in this example has also been previously reset.

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

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

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

If you then re-display the port status screen, you will see that the Intrusion Alert entry for port A3 has changed to “**No**”. That is, your evidence that the Intrusion Alert flag has been acknowledged (reset) is that the Intrusion Alert column in the port status display no longer shows “**Yes**” for the port on which the intrusion occurred (port A3 in this example). (Because the Intrusion Log

Configuring and Monitoring Port Security
Reading Intrusion Alerts and Resetting Alert Flags

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 9-15, 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 9-36.)

Syntax: show interfaces brief

List intrusion alert status (and other port status information).

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 A1.

```
ProCurve# show interfaces brief
```

Status and Counters - Port Status							
Port	Type	Intrusion			Status	Mode	Flow Ctrl
		Alert	Enabled				
A1	10/100TX	Yes	Yes	Up	10HDx	off	
A2	10/100TX	No	Yes	Up	10HDx	off	
A3	10/100TX	No	Yes	Up	10HDx	off	
A4	10/100TX	No	Yes	Up	10HDx	off	

Intrusion Alert on port A1.

Figure 9-16. 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:

```

ProCurve# show port-security intrusion-log
Status and Counters - Intrusion Log
Port  MAC Address          Date / Time
-----
A1    080009-e93d4f          07/03/02 21:09:34
A1    080009-21ae84          07/03/02 17:26:27
A1    080009-e93d4f prior to 07/03/02 17:18:43
    
```

MAC Address of latest Intruder on Port A1

Dates and Times of Intrusions

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

Figure 9-17. Example of the Intrusion Log with Multiple Entries for the Same Port

The above example shows three intrusions for port A1. Since the switch can show only one uncleared intrusion per port, the older two intrusions in this example have already been cleared by earlier use of the **clear intrusion-log** or the **port-security < 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 A1 and enable the switch to enter any subsequent intrusion for port A1 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 A1 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.)

```

ProCurve(config)# port-security a1 clear-intrusion-flag
ProCurve(config)# show interfaces brief
    
```

```

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

Intrusion Alert on port A1 is now cleared.

Figure 9-18. Example of Port Status Screen After Alert Flags Reset

For more on clearing intrusions, see “Note on Send-Disable Operation” on page 9-30.

Using the Event Log To Find Intrusion Alerts

The Event Log lists port security intrusions as:

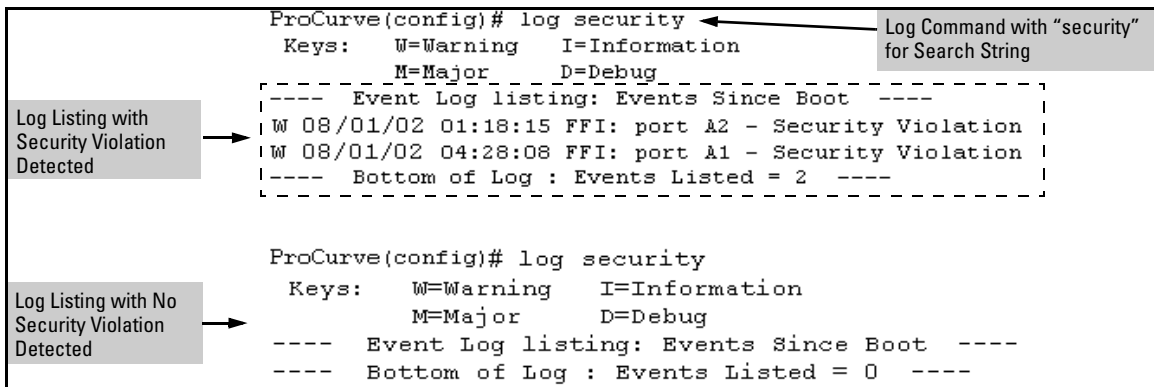
```
W MM/DD/YY HH:MM:SS FFI: port A3 - Security Violation
```

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



```
ProCurve(config)# log security ← Log Command with "security"
Keys:  W=Warning  I=Information
       M=Major    D=Debug
----- Event Log listing: Events Since Boot -----
| W 08/01/02 01:18:15 FFI: port A2 - Security Violation |
| W 08/01/02 04:28:08 FFI: port A1 - Security Violation |
|----- Bottom of Log : Events Listed = 2 -----|

ProCurve(config)# log security
Keys:  W=Warning  I=Information
       M=Major    D=Debug
----- Event Log listing: Events Since Boot -----
----- Bottom of Log : Events Listed = 0 -----
```

Log Listing with Security Violation Detected

Log Listing with No Security Violation Detected

Figure 9-19. 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. 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 chapter 11, “Using Authorized IP Managers”.)

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:

```
ProCurve(config)# port-security e a17 learn-mode static
address-limit 2
LACP has been disabled on secured port(s).
ProCurve(config)#
```

The switch will not allow you to configure LACP on a port on which port security is enabled. For example:

```
ProCurve(config)# int e a17 lacp passive
Error configuring port A17: LACP and port security cannot
be run together.
ProCurve(config)#
```

To restore LACP to the port, you must remove port security and re-enable LACP active or passive.

Traffic/Security Filters

Contents

Overview	10-2
Using Source-Port Filters	10-4
Operating Rules for Source-Port Filters	10-4
Configuring a Source-Port Filter	10-5
Viewing a Source-Port Filter	10-8
Filter Indexing	10-9
Editing a Source-Port Filter	10-10
Using Named Source-Port Filters	10-10

Overview

General Operation. You can enhance in-band security and improve control over access to network resources by configuring static per-port 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 between an inbound (source) port or trunk and any outbound (destination) ports and trunks (if any) on the switch. Source-port filters have no effect on traffic being routed across VLANs.

Note

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 10-5.

When you create a source port filter, all ports or port trunks on the switch appear as destinations on the list for that filter. The switch automatically forwards traffic to the 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.) For example, if you want 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 10-1 and 10-2).




Figure 10-1. Example of a Filter Blocking Traffic only from Port 5 to Server "A"

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
5	100/1000T	Forward
6	100/1000T	Forward
7	100/1000T	Drop
8	100/1000T	Forward
9	100/1000T	Forward
10	100/1000T	Forward
.	.	.
.	.	.
.	.	.
22	100/1000T	Forward
23	100/1000T	Forward
24	100/1000T	Forward

This list shows the filter created to block (drop) traffic from source port 5 (workstation "X") to destination port 7 (server "A"). Notice that the filter allows traffic to move from source port 5 to all other destination ports.

Figure 10-2. The Filter for the Actions Shown in Figure 10-1

Using Source-Port Filters

Operating Rules for Source-Port Filters

- You can configure one source-port filter for each physical port or port trunk on the switch.
- Each source-port filter you configure is composed of:
 - One source port or port trunk (**trk1**, **trk2**, ...**trk6**)
 - A set of destination ports and/or port trunks that includes all LAN ports and port trunks on the switch
 - An action 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.

Configuring a Source-Port Filter

The source-port filter command operates from the global configuration level.

Syntax: [no] filter source-port [e] < source-port-number > [drop | forward]

Creates or deletes the source port filter assigned to < source-port-number >. If you create a source-port filter without specifying a drop or forward action, the switch automatically creates a filter with a forward action from the designated source to all destinations on the switch.

[drop [e] < destination-port-list >]

Configures the filter for the designated source-port (or source-trunk) (< source-port-number >) to drop traffic for the ports and/or port trunks in the < destination-port-list >. Can be followed by the **forward** option if you have other destination ports set to **drop** that you want to change to **forward**. For example:

filter source-port <source-port-number> drop <destination-port-list > forward <destination-port-list >

[forward [e] < destination-port-list >]

*Configures the filter for the designated source (< source-port-number >) to forward traffic for the destinations 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 <source-port-number > forward < destination-port-list > drop < destination-port-list >

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:

```
ProCurve(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:

```
ProCurve(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 that 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:

```
ProCurve(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:

```
ProCurve(config)# filter source-port 5 drop 2
ProCurve(config)# trunk 5-6 trk1
ProCurve(config)# show filter
```

Traffic/Security Filters

IDX	Filter Type	Value
1	Source Port	*5*

```
ProCurve(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	Drop
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 10-3. Example of Switch Response to Adding a Filtered Source Port to a Trunk

Viewing a Source-Port Filter

You can list all source-port filters configured in the switch and, optionally, the detailed information on a specific filter.

Syntax: show filter [index]

Displays a listing of configured filters, where each filter entry includes an IDX (index) number, Filter Type, and Value :

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 (source-port) filter deletion created a gap in the filter listing.

Filter Type: Indicates the type of filter assigned to the IDX number.

Value: Indicates the port number or port-trunk name of the source port or trunk assigned to the filter.

Use **show filter** to learn the index number of a specific filter you want to examine in more detail.

[index]

*Displays detailed data on the filter designated by the index number. For source-port filters, the display includes the source-port number, a listing of all ports and/or trunks on the switch (with their port types), and the filter action configured on each port or trunk (**Forward**—the default— or **Drop**).*

For example, assume that these three filters exist on the switch:

Source Port	Destination Port(s)	Action
1	6-7	Drop; Forward on all other ports/trunks
2	8-9	Drop; Forward on all other ports/trunks
3	1-2	Drop; Forward on all other ports/trunks

If you wanted to determine the index number for the filter on source port 3 and then view a listing the filter details on source port 3, you would use the **show filter** and **show filter [index]** commands, as shown in figure 10-4.

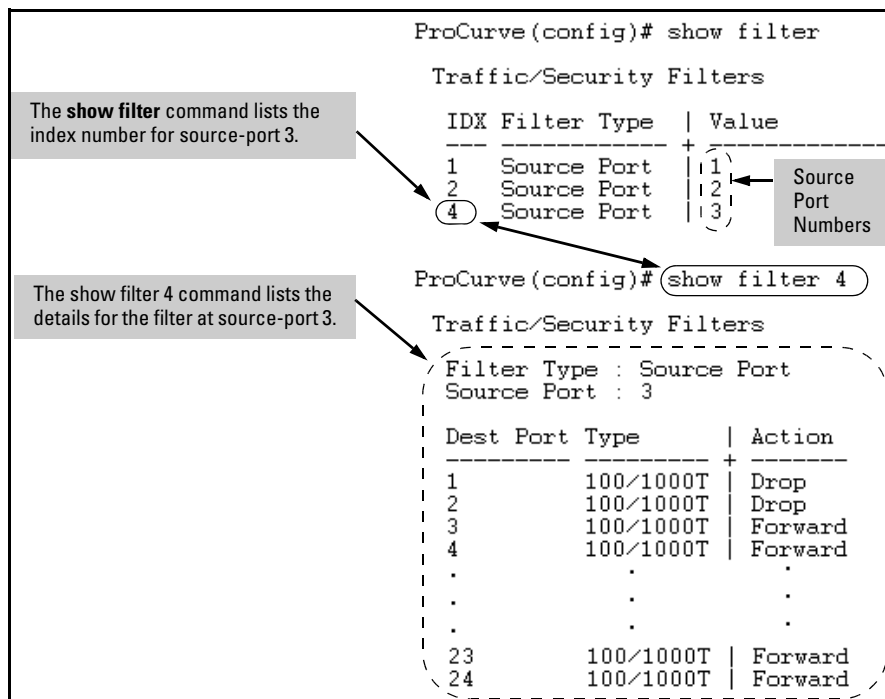


Figure 10-4. Example of Listing Filters and the Details of a Specific Filter

Filter Indexing

The switch automatically assigns each new source-port filter to the lowest-available index (IDX) number. 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.

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. 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 10-5. 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 10-5:

ProCurve(config)# show filter 1			ProCurve(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
8	100/1000T	Forward	8	100/1000T	Forward
9	100/1000T	Forward	9	100/1000T	Forward
10	100/1000T	Forward	10	100/1000T	Forward

Figure 10-5. Assigning Additional Destination Ports to an Existing Filter

Using Named Source-Port Filters

Named source-port filters are filters that may be used on multiple ports and port trunks. As with regular source-port filters, a port or port trunk can only have one source-port filter, but this new capability enables you to 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 < destination-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**.

```
ProCurve(config)# filter source-port named-filter web-only
ProCurve(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:

```
ProCurve(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.

Sample Configuration for 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 10-6. 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.

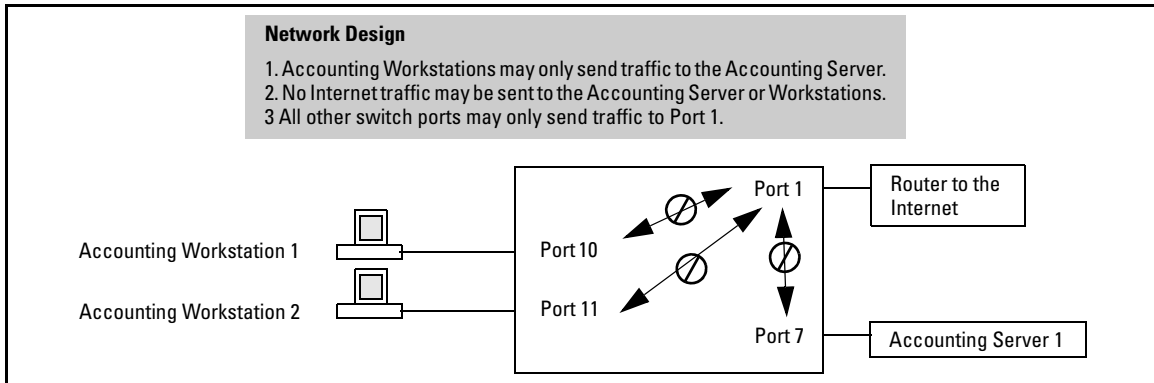


Figure 10-6. Network Configuration for Named Source-Port Filters Example

The company wants to use named source-port filters to direct inbound traffic only to the Internet while allowing only the two accounting workstations and the accounting server to communicate with each other, and not the Internet.

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.

Traffic/Security Filters

Using Source-Port Filters

```

ProCurve(config)# filter source-port named-filter web-only drop 2-26
ProCurve(config)# filter source-port named-filter accounting drop
1-6,8,9,12-26
ProCurve(config)# filter source-port named-filter no-incoming-web drop 7,10,11
ProCurve(config)#
ProCurve(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

```

ProCurve(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 10-7. Configuring and Defining Source Port Filters

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.

```

ProCurve(config)# filter source-port 2-6,8,9,12-26 named-filter web-only
ProCurve(config)# filter source-port 7,10,11 named-filter accounting
ProCurve(config)# filter source-port 1 named-filter no-incoming-web
ProCurve(config)#

```

Figure 10-8. Applying Source Port Filters to Ports

The **show filter** command shows what ports have filters applied.

```
ProCurve(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
9	Source Port	13
10	Source Port	14
11	Source Port	15
12	Source Port	16
13	Source Port	17
14	Source Port	18
15	Source Port	19
16	Source Port	20
17	Source Port	21
18	Source Port	22
19	Source Port	23
20	Source Port	24
21	Source Port	25
22	Source Port	26
23	Source Port	7

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 10-9. Show Filter with Source Ports

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.

ProCurve(config)# show filter 4			ProCurve(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	100/1000T	Forward	1	100/1000T	Drop
2	100/1000T	Forward	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	Drop
7	100/1000T	Forward	7	100/1000T	Forward
8	100/1000T	Forward	8	100/1000T	Drop
9	100/1000T	Forward	9	100/1000T	Drop
10	100/1000T	Forward	10	100/1000T	Forward
11	100/1000T	Forward	11	100/1000T	Forward
12	100/1000T	Forward	12	100/1000T	Drop

Figure 10-10. Show Filter on a Specific Port

The same command, using IDX 26, shows how traffic from the Internet is handled.

```
ProCurve(config)# show filter 26

Traffic/Security Filters

Filter Type : Source Port
Source Port : 1

Dest Port Type | Action
-----+-----
1      100/1000T | Forward
2      100/1000T | Forward
3      100/1000T | Forward
4      100/1000T | Forward
5      100/1000T | Forward
6      100/1000T | Forward
7      100/1000T | Drop
8      100/1000T | Forward
9      100/1000T | Forward
10     100/1000T | Drop
11     100/1000T | Drop
12     100/1000T | Forward
```

Figure 10-11. Show Filter for 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 port 8.

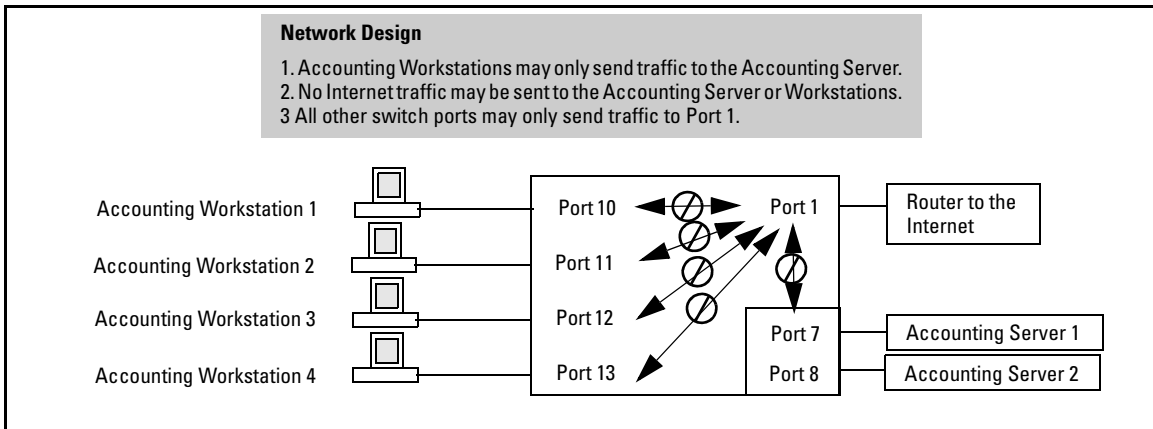


Figure 10-12. Expanded Network Configuration for Named Source-Port Filters Example

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.

Traffic/Security Filters

Using Source-Port Filters

```
ProCurve(config)# filter source-port named-filter accounting forward 8,12,13
ProCurve(config)# filter source-port named-filter no-incoming-web drop 8,12,13
ProCurve(config)#
ProCurve(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

ProCurve(config)#
```

Figure 10-13. Maintaining Traffic with Source Port Named-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.

```
ProCurve(config)# no filter source-port 8,12,13
ProCurve(config)# filter source-port 8,12,13 named-filter accounting
ProCurve(config)#
```

Figure 10-14. Applying Named Source-Port Filters to Switch Ports

The named source-port filters now manage traffic on the switch ports as shown below, using the **show filter source-port** command.

```
ProCurve(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
```

Figure 10-15. Managing Traffic with Named Source Port Filters

Using Authorized IP Managers

Contents

Overview	11-2
Configuration Options	11-3
Access Levels	11-3
Defining Authorized Management Stations	11-4
Overview of IP Mask Operation	11-4
Menu: Viewing and Configuring IP Authorized Managers	11-5
CLI: Viewing and Configuring Authorized IP Managers	11-6
Web: Configuring IP Authorized Managers	11-9
Building IP Masks	11-9
Configuring One Station Per Authorized Manager IP Entry	11-9
Configuring Multiple Stations Per Authorized Manager IP Entry ..	11-10
Additional Examples for Authorizing Multiple Stations	11-12
Operating Notes	11-12

Overview

Authorized IP Manager Features

Feature	Default	Menu	CLI	Web
Listing (Showing) Authorized Managers	n/a	page 11-5	page 11-6	page 11-9
Configuring Authorized IP Managers	None	page 11-5	page 11-6	page 11-9
Building IP Masks	n/a	page 11-9	page 11-9	page 11-9
Operating and Troubleshooting Notes	n/a	page 11-12	page 11-12	page 11-12

The Authorized IP Managers feature uses IP addresses and masks to determine 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
- SNMP (with a correct community name)

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 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. Refer to table 1-1, “Management Access Security Protection” (page 1-4) for a listing of access security features with the security coverage they provide.

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 (for Telnet, SNMPv1, and SNMPv2c access only)

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 authorized station "spoofs" an authorized IP address, it can gain management 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

The Authorized IP Manager feature can assign an access level to stations using Telnet, SNMPv1, or SNMPv2c for switch access. The access level the switch allows for authorized stations using SSH, SNMPv3, or the web browser interface is determined by the access application itself, and not by the Authorized IP Manager feature. The IP Authorized list does not enforce access rights because SSH, the Web Agent (SSL) and SNMPv3 have much better access control mechanisms than an IP address alone. For example, SNMPv3 has access control down to the user level as well as authentication and encryption for data integrity.

For each authorized manager address using Telnet, SNMPv1, or SNMPv2c, you can configure either of these access levels:

- **Manager:** Enables full access to all web browser and console interface screens for viewing, configuration, and all other operations available in these interfaces.

- **Operator:** Allows read-only access from the web browser and console interfaces. (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 management station to have IP access to the switch. To use this method, just enter the IP address of an authorized management station in the Authorized 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 11-9.)
- **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 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 11-10.)

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 management station.

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 Authorized 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 11-9.

Note

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

From the console Main Menu, select:

2. Switch Configuration ...

7. IP Authorized Managers

```
----- TELNET - MANAGER MODE -----
Switch Configuration - IP Managers

Authorized Manager IP      IP Mask      Access Level
-----
13.28.227.101             255.255.255.252  Manager
13.28.227.104             255.255.255.254  Manager
13.28.227.106             255.255.255.0    Operator
13.28.227.125             255.255.255.255  Manager

Actions->  Back  Add  Edit  Delete  Help

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

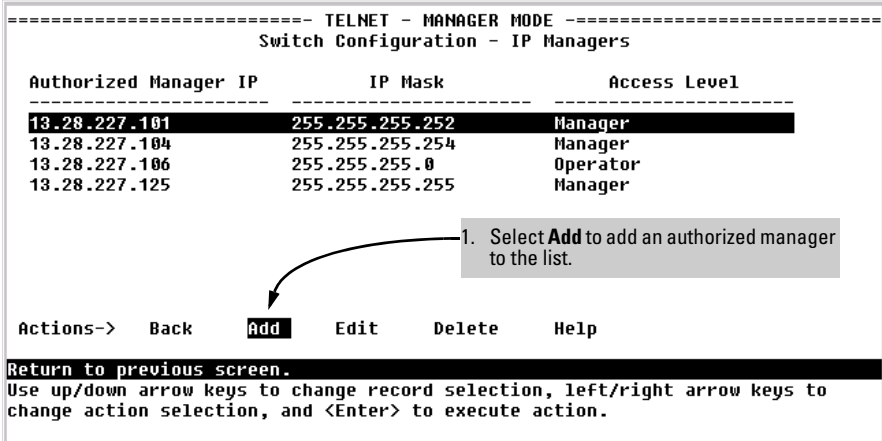


Figure 11-1. Example of How To Add an Authorized Manager Entry

Using Authorized IP Managers

Defining Authorized Management Stations

```

----- CONSOLE - MANAGER MODE -----
Switch Configuration - IP

Authorized Manager IP : ██████████
IP Mask [255.255.255.255] : 255.255.255.255
Access Level : Manager

Actions->  _Cancel    _Edit    _Save    _Help

Enter the IP address of an authorized manager.
Use arrow keys to change field selection. <Space>
and <Enter> to go to Actions.
  
```

2. Enter an Authorized Manager IP address here.

3. Use the default mask to allow access by one management device, or edit the mask to allow access by a block of management devices. See "Building IP Masks" on page 11-9.

4. Use the Space bar to select Manager or Operator access.

5. Press [Enter], then [S] (for Save) to configure the IP Authorized Manager entry.

Applies only to access through Telnet, SNMPv1, and SNMPv2c.

Figure 11-2. Example of How To Add an Authorized Manager Entry (Continued)

Editing or Deleting an Authorized Manager Entry. Go to the IP Managers List screen (figure 11-1), 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	11-7
<ip-address>	11-8
<ip-mask-bits>	11-8
[access <operator manager>]	

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:

```

ProCurve# show ip authorized-managers
IP Managers
Authorized Manager IP   IP Mask                               Access Level
-----
10.28.227.101          255.255.255.252                     Manager
10.28.227.104          255.255.255.254                     Manager
10.28.227.125          255.255.255.255                     Manager
10.28.227.106          255.255.255.0                       Operator
  
```

Figure 11-3. Example of the Show IP Authorized-Manager Display

The above example shows an Authorized IP Manager List that allows stations to access the switch as shown below:

IP Mask	Authorized Station IP Address:	Access Mode:
255.255.255.252	10.28.227.100 through 103	Manager
255.255.255.254	10.28.227.104 through 105	Manager
255.255.255.255	10.28.227.125	Manager
255.255.255.0	10.28.227.0 through 255	Operator

Configuring IP Authorized Managers for the Switch

Syntax: ip authorized-managers < ip address >

Configures one or more authorized IP addresses.

[< ip-mask-bits >]

Configures the IP mask for < ip address >

[access < operator | manager >]

Configures the privilege level for < ip address >.

Applies only to access through Telnet, SNMPv1, and SNMPv2c. Refer to the Note on page 11-3.

To Authorize Manager Access. This command authorizes manager-level access for any station having an IP address of 10.28.227.0 through 10.28.227.255:

```

ProCurve(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:

```
ProCurve(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** for the mask. If you do not specify either Manager or Operator access, the switch automatically assigns the Manager access. For example:

```
ProCurve Switch 2824(config)# ip authorized-managers 10.28.227.105  
ProCurve Switch 2824(config)# show ip authorized-managers
```

Authorized Manager IP	IP Mask	Access Level
10.28.227.105	255.255.255.255	Manager

Omitting a mask in the ip authorized-managers command results in a default mask of 255.255.255.255, which authorizes only the specified station. Refer to "Configuring Multiple Stations Per Authorized Manager IP Entry" on page 11-10.

Figure 11-4. Example of Specifying an IP Authorized Manager with the Default Mask

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.):

```
ProCurve(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.

```
ProCurve(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:

```
ProCurve(config)# no ip authorized-managers 10.28.227.101
```


Web: Configuring IP Authorized Managers

In the web browser interface you can configure IP Authorized Managers as described below.

To Add, Modify, or Delete an IP Authorized Manager address:

1. Click on the **Security** tab.
2. Click on **[Authorized Addresses]**.
3. Enter the appropriate parameter settings for the operation you want.
4. Click on **[Add]**, **[Replace]**, or **[Delete]** to implement the configuration change.

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 quickly by simply 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 11-3 on page 11-7, if you configure an IP address of **10.28.227.125** with an IP mask of **255.255.255.255**, only a station with an IP address of **10.28.227.125** has management access to the switch.

Table 11-1. 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	28	227	125	

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 11-2. 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	

	1st Octet	2nd Octet	3rd Octet	4th Octet	Manager-Level or Operator-Level Device Access
IP Mask	255	255	255	249	In this example (figure 11-5, 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	

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)	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	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. 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: <ul style="list-style-type: none"> • 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.
4th Octet of IP Authorized Address (125)	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	

Figure 11-5. Example of How the Bitmap in the IP Mask Defines Authorized Manager Addresses

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.

—This page is intentionally unused—

Index

Numerics

3DES ... 6-3, 7-3

802.1X

See *port-based access control*. ... 8-1

A

aaa authentication ... 4-8

aaa port-access

See *Web or MAC Authentication*.

access levels, authorized IP managers ... 11-3

accounting

See *RADIUS*.

address

authorized for port security ... 9-3

authentication

See *TACACS*.

authorized addresses

for IP management security ... 11-4

for port security ... 9-3

authorized IP managers

access levels ... 11-3

building IP masks ... 11-9

configuring in browser interface ... 11-7, 11-9

configuring in console ... 11-5

definitions of single and multiple ... 11-4

effect of duplicate IP addresses ... 11-12

IP mask for multiple stations ... 11-10

IP mask for single station ... 11-9

IP mask operation ... 11-4

operating notes ... 11-12

overview ... 11-1

precedence over other security ... 11-2

troubleshooting ... 11-12

C

certificate

CA-signed ... 7-4

root ... 7-4

self-signed ... 7-4

Clear button

to delete password protection ... 2-5

configuration

port security ... 9-5

RADIUS

See *RADIUS*.

SSH

See *SSH*.

connection inactivity time ... 2-3

console, for configuring

authorized IP managers ... 11-5

D

DES ... 6-3, 7-3

disclaimer ... 1-ii

duplicate IP address

effect on authorized IP managers ... 11-12

E

event log

intrusion alerts ... 9-35

F

filter, source-port

configuring ... 10-4

editing ... 10-9

filter indexing ... 10-8

filter type ... 10-7

idx ... 10-7, 10-8

index ... 10-7, 10-8

named source-port filters ... 10-9

operating rules ... 10-3, 10-10

operation ... 10-2

port-trunk operation ... 10-2, 10-5

show ... 10-7, 10-12

value ... 10-7

viewing ... 10-7, 10-12

G

GVRP, static VLAN not advertised ... 8-47

I

- inconsistent value, message ... 9-14
- intrusion alarms
 - entries dropped from log ... 9-36
 - event log ... 9-35
 - prior to ... 9-36
- Intrusion Log
 - prior to ... 9-31, 9-34
- IP
 - authorized IP managers ... 11-1
 - reserved port numbers ... 6-17
- IP masks
 - building ... 11-9
 - for multiple authorized manager stations ... 11-10
 - for single authorized manager station ... 11-9
 - operation ... 11-4

K

- kill command ... 6-11

L

- LACP
 - 802.1X not allowed ... 8-11, 8-15, 8-48

M

- MAC Authentication
 - authenticator operation ... 3-5
 - blocked traffic ... 3-4
 - CHAP
 - defined ... 3-9
 - usage ... 3-4
 - client status ... 3-29
 - configuration commands ... 3-23
 - configuring
 - on the switch ... 3-22
 - switch for RADIUS access ... 3-15
 - the RADIUS server ... 3-14
 - features ... 3-4
 - general setup ... 3-12
 - LACP not allowed ... 3-11
 - rules of operation ... 3-10
 - show status and configuration ... 3-27
 - terminology ... 3-9
- manager password ... 2-2, 2-4

- manager password recommended ... 4-7
- MD5

See *RADIUS*.

- message
 - inconsistent value ... 9-14

O

- open VLAN mode
 - See *port access control*
- OpenSSH ... 6-3
- OpenSSL ... 7-2
- operating notes
 - authorized IP managers ... 11-12
 - port security ... 9-36
- operator password ... 2-2, 2-4

P

- password
 - authorized IP managers, precedence ... 11-2
 - browser/console access ... 2-3
 - case-sensitive ... 2-4
 - caution ... 2-3
 - delete ... 2-4
 - deleting with the Clear button ... 2-5
 - if you lose the password ... 2-5
 - incorrect ... 2-3
 - length ... 2-4
 - operator only, caution ... 2-3
 - pair ... 2-2
 - setting ... 2-4
- password pair ... 2-2
- password security ... 6-18
- port
 - security configuration ... 9-2
- port security
 - authorized address definition ... 9-3
 - authorized IP managers, precedence ... 11-2
 - basic operation ... 9-2
 - configuring ... 9-5
 - configuring in browser interface ... 9-27, 9-36
 - event log ... 9-35
 - notice of security violations ... 9-28
 - operating notes ... 9-36
 - overview ... 9-2
 - prior to ... 9-36
 - proxy web server ... 9-36

- port-based access control
 - authenticate switch ... 8-4
 - authenticate users ... 8-4
 - authenticator backend state ... 8-38
 - authenticator operation ... 8-6, 8-8
 - authenticator, show commands ... 8-38
 - authorized IP managers, precedence ... 11-2
 - block traffic ... 8-3
 - blocking non-802.1X device ... 8-33
 - CHAP ... 8-3
 - chap-radius ... 8-19
 - configuration commands ... 8-15
 - configuration overview ... 8-13
 - configuration, displaying ... 8-38
 - configuring method ... 8-19
 - counters ... 8-38
 - EAP ... 8-3
 - EAPOL ... 8-9
 - eap-radius ... 8-19
 - enabling on ports ... 8-15
 - enabling on switch ... 8-20
 - features ... 8-3
 - general setup ... 8-12
 - GVRP effect ... 8-47
 - LACP not allowed ... 8-48
 - local ... 8-19
 - local username and password ... 8-4
 - MD5 ... 8-8
 - messages ... 8-48
 - open VLAN
 - authorized client ... 8-22
 - configuration ... 8-28, 8-30
 - general operation ... 8-21
 - mode ... 8-21
 - operating notes ... 8-31
 - operating rules ... 8-25
 - PVID, no ... 8-40
 - security breach ... 8-31
 - set up ... 8-27
 - status, viewing ... 8-40
 - suspended VLAN ... 8-41
 - unauthorized client ... 8-22
 - use models ... 8-22
 - VLAN, after authentication ... 8-22, 8-26, 8-31
 - VLAN, tagged ... 8-21, 8-22, 8-23, 8-26, 8-31, 8-42
 - operation ... 8-6
 - overview ... 8-3
 - port-security, with 802.1X ... 8-32
 - RADIUS ... 8-3
 - RADIUS host IP address ... 8-20
 - rules of operation ... 8-10
 - show commands ... 8-38
 - show commands, supplicant ... 8-43
 - statistics ... 8-38
 - supplicant operation ... 8-8
 - supplicant operation, switch-port ... 8-7
 - supplicant state ... 8-43
 - supplicant statistics, note ... 8-43
 - supplicant, configuring ... 8-34
 - supplicant, configuring switch port ... 8-36
 - supplicant, enabling ... 8-35
 - switch username and password ... 8-4
 - terminology ... 8-8
 - troubleshooting, gvrp ... 8-44
 - used with port-security ... 8-32
 - VLAN operation ... 8-44
 - prior to ... 9-31, 9-34, 9-36
 - Privacy Enhanced Mode (PEM)
 - See *SSH*.
 - proxy
 - web server ... 9-36
- ## Q
- quick start ... 1-8
- ## R
- ### RADIUS
- accounting ... 5-2, 5-17
 - accounting, configuration outline ... 5-19
 - accounting, configure server access ... 5-20
 - accounting, configure types on switch ... 5-22
 - accounting, exec ... 5-18, 5-22
 - accounting, interim updating ... 5-24
 - accounting, network ... 5-22
 - accounting, operating rules ... 5-19
 - accounting, server failure ... 5-19
 - accounting, session-blocking ... 5-24
 - accounting, start-stop method ... 5-23
 - accounting, statistics terms ... 5-26
 - accounting, stop-only method ... 5-23
 - accounting, system ... 5-18, 5-22
 - authentication options ... 5-2

- authentication, local ... 5-16
- authorized IP managers, precedence ... 11-2
- bypass RADIUS server ... 5-9
- commands, accounting ... 5-17
- commands, switch ... 5-6
- configuration outline ... 5-7
- configure server access ... 5-10
- configuring switch global parameters ... 5-12
- general setup ... 5-5
- local authentication ... 5-9
- MD5 ... 5-4
- messages ... 5-32
- network accounting ... 5-18
- operating rules, switch ... 5-4
- security ... 5-9
- security note ... 5-2
- server access order ... 5-19
- server access order, changing ... 5-30
- servers, multiple ... 5-13
- show accounting ... 5-28
- show authentication ... 5-27
- SNMP access security not supported ... 5-2
- statistics, viewing ... 5-25
- terminology ... 5-3
- TLS ... 5-4
- Web browser authentication ... 5-7
- web-browser access controls ... 5-17
- web-browser security not supported ... 5-2, 5-17
- RADIUS accounting
 - See *RADIUS*.
- reserved port numbers ... 6-17, 7-20

S

- security
 - authorized IP managers ... 11-1
 - per port ... 9-2
- security violations
 - notices of ... 9-28
- security, password
 - See *SSH*.
- setting a password ... 2-4
- setup screen ... 1-8
- show
 - locked down MAC addresses ... 9-25
 - locked out MAC addresses ... 9-26
- SSH
 - authenticating switch to client ... 6-3

- authentication, client public key ... 6-2
- authentication, user password ... 6-2
- caution, restricting access ... 6-20
- caution, security ... 6-18
- CLI commands ... 6-9
- client behavior ... 6-15, 6-16
- client public-key authentication ... 6-19, 6-22
- client public-key, clearing ... 6-27
- client public-key, creating file ... 6-24
- client public-key, displaying ... 6-26
- configuring authentication ... 6-18
- crypto key ... 6-11
- disabling ... 6-11
- enable ... 6-16, 7-19
- enabling ... 6-15
- erase host key pair ... 6-11
- generate host key pair ... 6-11
- generating key pairs ... 6-10
- host key pair ... 6-11
- key, babble ... 6-11
- key, fingerprint ... 6-11
- keys, zeroing ... 6-11
- key-size ... 6-17
- known-host file ... 6-13, 6-15
- man-in-the-middle spoofing ... 6-16
- messages, operating ... 6-28
- OpenSSH ... 6-3
- operating rules ... 6-8
- outbound SSH not secure ... 6-8
- password security ... 6-18
- password-only authentication ... 6-18
- passwords, assigning ... 6-9
- PEM ... 6-4
- prerequisites ... 6-5
- public key ... 6-5, 6-13
- public key, displaying ... 6-14
- reserved IP port numbers ... 6-17
- security ... 6-18
- SSHv2 ... 6-2
- stacking, security ... 6-8
- steps for configuring ... 6-6
- supported encryption methods ... 6-3
- switch key to client ... 6-13
- terminology ... 6-4
- unauthorized access ... 6-20, 6-28
- version ... 6-2
- zeroing a key ... 6-11
- zeroize ... 6-11

SSL

- CA-signed ... 7-4, 7-15
 - CA-signed certificate ... 7-4, 7-15
 - CLI commands ... 7-7
 - client behavior ... 7-17, 7-18
 - crypto key ... 7-10
 - disabling ... 7-10
 - enabling ... 7-17
 - erase certificate key pair ... 7-10
 - erase host key pair ... 7-10
 - generate CA-signed certificate ... 7-15
 - generate host key pair ... 7-10
 - generate self-signed ... 7-13
 - generate self-signed certificate ... 7-10, 7-13
 - generate server host certificate ... 7-10
 - generating Host Certificate ... 7-9
 - host key pair ... 7-10
 - key, babble ... 7-12
 - key, fingerprint ... 7-12
 - man-in-the-middle spoofing ... 7-18
 - OpenSSL ... 7-2
 - operating notes ... 7-6
 - operating rules ... 7-6
 - passwords, assigning ... 7-7
 - prerequisites ... 7-5
 - remove self-signed certificate ... 7-10
 - remove server host certificate ... 7-10
 - reserved TCP port numbers ... 7-20
 - root ... 7-4
 - root certificate ... 7-4
 - self-signed ... 7-4, 7-13
 - self-signed certificate ... 7-4, 7-10, 7-13
 - server host certificate ... 7-10
 - SSL server ... 7-3
 - SSLv3 ... 7-2
 - stacking, security ... 7-6
 - steps for configuring ... 7-5
 - supported encryption methods ... 7-3
 - terminology ... 7-3
 - TLSv1 ... 7-2
 - troubleshooting, operating ... 7-21
 - version ... 7-2
 - zeroize ... 7-10, 7-12
- stacking
- SSH security ... 6-8
 - SSL security ... 7-6

T

TACACS

- aaa parameters ... 4-12
 - authentication ... 4-3
 - authentication process ... 4-20
 - authentication, local ... 4-22
 - authorized IP managers, effect ... 4-26
 - authorized IP managers, precedence ... 11-2
 - configuration, authentication ... 4-11
 - configuration, encryption key ... 4-19
 - configuration, server access ... 4-15
 - configuration, timeout ... 4-20
 - configuration, viewing ... 4-10
 - encryption key ... 4-6, 4-15, 4-16, 4-19
 - encryption key, general operation ... 4-23
 - encryption key, global ... 4-20
 - general operation ... 4-2
 - IP address, server ... 4-15
 - local manager password requirement ... 4-26
 - messages ... 4-25
 - NAS ... 4-3
 - overview ... 1-2
 - precautions ... 4-5
 - preparing to configure ... 4-8
 - preventing switch lockout ... 4-15
 - privilege level code ... 4-7
 - server access ... 4-15
 - server priority ... 4-18
 - setup, general ... 4-5
 - show authentication ... 4-8
 - system requirements ... 4-5
 - TACACS+ server ... 4-3
 - testing ... 4-5
 - timeout ... 4-15
 - troubleshooting ... 4-6
 - unauthorized access, preventing ... 4-7
 - web access, controlling ... 4-24
 - web access, no effect on ... 4-5
- tacacs-server ... 4-8
- ### TCP
- reserved port numbers ... 7-20
- ### TLS
- See *RADIUS*.
- ### troubleshooting
- authorized IP managers ... 11-12
- ### trunk
- filter, source-port ... 10-2, 10-5
 - LACP, 802.1X not allowed ... 8-15

See also *LACP*.

U

user name
cleared ... 2-5

V

value, inconsistent ... 9-14
VLAN
802.1X ... 8-44
802.1X, ID changes ... 8-47
802.1X, suspend untagged VLAN ... 8-41
not advertised for GVRP ... 8-47

W

warranty ... ii
Web Authentication
authenticator operation ... 3-5
blocked traffic ... 3-4
CHAP
defined ... 3-9
usage ... 3-4
client status ... 3-29
configuration commands ... 3-18
configuring
on the switch ... 3-17
switch for RADIUS access ... 3-15
features ... 3-4
general setup ... 3-12
LACP not allowed ... 3-11
redirect URL ... 3-9
rules of operation ... 3-10
show status and configuration ... 3-26
terminology ... 3-9
web browser interface, for configuring
authorized IP managers ... 11-7, 11-9
web browser interface, for configuring port
security ... 9-27, 9-36
web server, proxy ... 9-36

— *This page is intentionally unused.* —



Technical information in this document is subject to change without notice.

© Copyright Hewlett-Packard Company, 2000, 2004, 2005, 2006. All rights reserved. Reproduction, adaptation, or translation without prior written permission is prohibited except as allowed under the copyright laws.

July 2006

Manual Part Number
5991-4734