

Access Security Guide

2610
2610-PWR

ProCurve Switches
R.11.XX

www.procurve.com



ProCurve
Switch 2610 Series
Switch 2610-PWR Series

December 2007

Access Security Guide

© Copyright 2007 Hewlett-Packard Company, L.P.
The information contained herein is subject to change without notice.

Publication Number

5991-8642
December 2007

Applicable Products

ProCurve Switch 2610-24	(J9085A)
ProCurve Switch 2610-48	(J9088A)
ProCurve Switch 2610-24-PWR	(J9087A)
ProCurve Switch 2610-48-PWR	(J9089A)
ProCurve Switch 2610-24/12-PWR	(J9086A)

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

Software Feature Index	xiv
------------------------------	-----

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	2-16
Password Recovery Process	2-18

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	3-13
Configuring the Switch To Access a RADIUS Server	3-15
Configuring Web Authentication	3-18
Overview	3-18
Configure the Switch for Web-Based Authentication	3-19
Configuring MAC Authentication on the Switch	3-23
Overview	3-23
Configure the Switch for MAC-Based Authentication	3-24
Show Commands for Web-Based Authentication	3-28
Show Commands for MAC-Based Authentication	3-31
Show Client Status	3-33

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	4-10
Configuring the Switch's TACACS+ Server Access	4-17
How Authentication Operates	4-22
General Authentication Process Using a TACACS+ Server	4-22
Local Authentication Process	4-24
Using the Encryption Key	4-25
Controlling Web Browser Interface Access When Using TACACS+ Authentication	4-26
Messages Related to TACACS+ Operation	4-27
Operating Notes	4-28

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	5-11
3. Configure the Switch's Global RADIUS Parameters	5-13
Local Authentication Process	5-17

Controlling Web Browser Interface Access When Using RADIUS Authentication	5-18
Configuring RADIUS Authorization	5-18
Overview	5-18
Commands Authorization Type	5-19
Enabling Authorization with the CLI	5-19
Showing Authorization Information	5-20
Configuring the RADIUS Server	5-20
Configuring RADIUS Accounting	5-26
Operating Rules for RADIUS Accounting	5-27
Steps for Configuring RADIUS Accounting	5-28
Viewing RADIUS Statistics	5-33
General RADIUS Statistics	5-33
RADIUS Authentication Statistics	5-35
RADIUS Accounting Statistics	5-36
Changing RADIUS-Server Access Order	5-37
Messages Related to RADIUS Operation	5-39

6 Configuring RADIUS Server Support for Switch Services

Contents	6-1
Overview	6-2
Configuring the RADIUS Server for CoS Services	6-3
Viewing the Currently Active Per-Port CoS Configuration Specified by a RADIUS Server	6-3
Configuring and Using RADIUS-Assigned Access Control Lists	6-6
Introduction	6-6
Terminology	6-6
Overview of RADIUS-Assigned, Dynamic Port ACLs	6-9
Contrasting Dynamic and Static ACLs	6-11
How a RADIUS Server Applies a Dynamic Port ACL to a Switch Port	6-12
General ACL Features, Planning, and Configuration	6-13

The Packet-filtering Process	6-14
Operating Rules for Dynamic Port ACLs	6-14
Configuring an ACL in a RADIUS Server	6-15
Configuring ACE Syntax in RADIUS Servers	6-18
Configuring the Switch To Support Dynamic Port ACLs	6-20
Displaying the Current Dynamic Port ACL Activity on the Switch	6-21
Event Log Messages	6-24
Causes of Client Deauthentication Immediately After Authenticating	6-25
Monitoring Shared Resources	6-25

7 Configuring Secure Shell (SSH)

Contents	7-1
Overview	7-2
Terminology	7-4
Prerequisite for Using SSH	7-5
Public Key Formats	7-5
Steps for Configuring and Using SSH for Switch and Client Authentication	7-6
General Operating Rules and Notes	7-8
Configuring the Switch for SSH Operation	7-9
1. Assign Local Login (Operator) and Enable (Manager) Password ..	7-9
2. Generate the Switch's Public and Private Key Pair	7-10
3. Provide the Switch's Public Key to Clients	7-12
4. Enable SSH on the Switch and Anticipate SSH Client Contact Behavior	7-15
5. Configure the Switch for SSH Authentication	7-18
6. Use an SSH Client To Access the Switch	7-21
Further Information on SSH Client Public-Key Authentication	7-22
Messages Related to SSH Operation	7-28

8 Configuring Secure Socket Layer (SSL)

Contents	8-1
Overview	8-2
Terminology	8-3
Prerequisite for Using SSL	8-5
Steps for Configuring and Using SSL for Switch and Client Authentication	8-5
General Operating Rules and Notes	8-6
Configuring the Switch for SSL Operation	8-7
1. Assign Local Login (Operator) and Enable (Manager) Password .	8-7
2. Generate the Switch's Server Host Certificate	8-9
3. Enable SSL on the Switch and Anticipate SSL Browser Contact Behavior	8-17
Common Errors in SSL Setup	8-21

9 Access Control Lists (ACLs)

Contents	9-1
Introduction	9-3
ACL Applications	9-3
Optional Network Management Applications	9-3
Optional PCM and IDM Applications	9-4
General Application Options	9-4
Terminology	9-6
Overview	9-9
Types of IP ACLs	9-9
ACL Inbound Application Points	9-9
Features Common to All ACLs	9-10
General Steps for Planning and Configuring ACLs	9-11
ACL Operation	9-12
Introduction	9-12
The Packet-Filtering Process	9-13
Planning an ACL Application	9-16
Switch Resource Usage	9-16

Managing ACL Resource Consumption	9-18
Traffic Management and Improved Network Performance	9-22
Security	9-22
Guidelines for Planning the Structure of an ACL	9-23
ACL Configuration and Operating Rules	9-24
How an ACE Uses a Mask To Screen Packets for Matches	9-25
Configuring and Assigning an ACL	9-32
Overview	9-32
ACL Configuration Structure	9-33
ACL Configuration Factors	9-36
Using the CLI To Create an ACL	9-38
Configuring and Assigning a Numbered, Standard ACL	9-39
Configuring and Assigning a Numbered, Extended ACL	9-44
Configuring a Named ACL	9-50
Enabling or Disabling ACL Filtering on an Interface	9-52
Deleting an ACL from the Switch	9-53
Displaying ACL Data	9-54
Display an ACL Summary	9-54
Display the Content of All ACLs on the Switch	9-55
Display the ACL Assignments for an Interface	9-56
Displaying the Content of a Specific ACL	9-57
Displaying the Current ACL Resources	9-59
Display All ACLs and Their Assignments in the Switch Startup-Config File and Running-Config File	9-60
Editing ACLs and Creating an ACL Offline	9-60
Using the CLI To Edit ACLs	9-60
Working Offline To Create or Edit an ACL	9-63
Enable ACL “Deny” Logging	9-67
Requirements for Using ACL Logging	9-67
ACL Logging Operation	9-67
Enabling ACL Logging on the Switch	9-68
Operating Notes for ACL Logging	9-70
General ACL Operating Notes	9-71

10 Traffic/Security Filters

Contents	10-1
Overview	10-2
General Operation	10-2
Applying a Source Port Filter in a Multinetted VLAN	10-3
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	10-7
Filter Indexing	10-9
Editing a Source-Port Filter	10-9
Using Named Source-Port Filters	10-10

11 Configuring Port-Based and User-Based Access Control (802.1X)

Contents	11-1
Overview	11-3
Why Use Port-Based or User-Based Access Control?	11-3
General Features	11-3
User Authentication Methods	11-4
Terminology	11-6
General 802.1X Authenticator Operation	11-9
Example of the Authentication Process	11-9
VLAN Membership Priority	11-10
General Operating Rules and Notes	11-12
General Setup Procedure for 802.1X Access Control	11-14
Do These Steps Before You Configure 802.1X Operation	11-14
Overview: Configuring 802.1X Authentication on the Switch	11-16
Configuring Switch Ports as 802.1X Authenticators	11-17
1. Enable 802.1X Authentication on Selected Ports	11-18
2. Reconfigure Settings for Port-Access	11-20
3. Configure the 802.1X Authentication Method	11-24
4. Enter the RADIUS Host IP Address(es)	11-25

5. Enable 802.1X Authentication on the Switch	11-26
6. Optional: Reset Authenticator Operation	11-26
7. Optional: Configure 802.1X Controlled Directions	11-26
802.1X Open VLAN Mode	11-29
Introduction	11-29
VLAN Membership Priorities	11-30
Use Models for 802.1X Open VLAN Modes	11-31
Operating Rules for Authorized-Client and Unauthorized-Client VLANs	11-36
Setting Up and Configuring 802.1X Open VLAN Mode	11-40
802.1X Open VLAN Operating Notes	11-44
Option For Authenticator Ports: Configure Port-Security To Allow Only 802.1X-Authenticated Devices	11-45
Port-Security	11-46
Configuring Switch Ports To Operate As Supplicants for 802.1X Connections to Other Switches	11-47
Example	11-47
Supplicant Port Configuration	11-48
Displaying 802.1X Configuration, Statistics, and Counters	11-51
Show Commands for Port-Access Authenticator	11-51
Viewing 802.1X Open VLAN Mode Status	11-54
Show Commands for Port-Access Supplicant	11-57
How RADIUS/802.1X Authentication Affects VLAN Operation	11-58
VLAN Assignment on a Port	11-59
Operating Notes	11-59
Example of Untagged VLAN Assignment in a RADIUS-Based Authentication Session	11-61
Enabling the Use of GVRP-Learned Dynamic VLANs in Authentication Sessions	11-64
Operating Note	11-66
Messages Related to 802.1X Operation	11-67

12 Configuring and Monitoring Port Security

Contents	12-1
Overview	12-2
Basic Operation	12-2
Eavesdrop Protection	12-3
Blocking Unauthorized Traffic	12-3
Trunk Group Exclusion	12-4
Planning Port Security	12-5
Port Security Command Options and Operation	12-6
Retention of Static MAC Addresses	12-10
Displaying Current Port Security Settings	12-10
Configuring Port Security	12-12
MAC Lockdown	12-17
Differences Between MAC Lockdown and Port Security	12-19
Deploying MAC Lockdown	12-21
MAC Lockout	12-25
Port Security and MAC Lockout	12-27
Web: Displaying and Configuring Port Security Features	12-27
Reading Intrusion Alerts and Resetting Alert Flags	12-28
Notice of Security Violations	12-28
How the Intrusion Log Operates	12-29
Keeping the Intrusion Log Current by Resetting Alert Flags	12-29
Using the Event Log To Find Intrusion Alerts	12-34
Web: Checking for Intrusions, Listing Intrusion Alerts, and Resetting Alert Flags	12-35
Operating Notes for Port Security	12-35

13 Using Authorized IP Managers

Contents	13-1
Overview	13-2
Configuration Options	13-3
Access Levels	13-3
Defining Authorized Management Stations	13-4

Overview of IP Mask Operation	13-4
Menu: Viewing and Configuring IP Authorized Managers	13-5
CLI: Viewing and Configuring Authorized IP Managers	13-6
Web: Configuring IP Authorized Managers	13-9
Web Proxy Servers	13-9
Web-Based Help	13-10
Building IP Masks	13-10
Configuring One Station Per Authorized Manager IP Entry	13-10
Configuring Multiple Stations Per Authorized Manager IP Entry ..	13-11
Additional Examples for Authorizing Multiple Stations	13-13
Operating Notes	13-13

Index

Product Documentation

Note

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

Printed Publications

The two publications listed below are printed and shipped with your switch. The latest version of each is also available in PDF format on the ProCurve Web site, as described in the Note at the top of this page.

- *Read Me First*—Provides software update information, product notes, and other information.
- *Installation and Getting Started Guide*—Explains how to prepare for and perform the physical installation and connect the switch to your network.

Electronic Publications

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

- *Management and Configuration Guide*—Describes how to configure, manage, and monitor basic switch operation.
- *Advanced Traffic Management Guide*—Explains how to configure traffic management features, such as spanning tree, VLANs, and IP routing.
- *Access Security Guide*—Explains how to configure access security features and user authentication on the switch.
- *Release Notes*—Describe new features, fixes, and enhancements that become available between revisions of the above guides.

Software Feature Index

For the software manual set supporting your switch model, the following feature index indicates which manual to consult for information on a given software feature. (Note that some software features are not supported on all switch models.)

Feature	Management and Configuration	Advanced Traffic Management	Access Security Guide
802.1Q VLAN Tagging	-	X	-
802.1X Port-Based Priority	X	-	-
ACLs	-	-	X
AAA Authentication	-	-	X
Authorized IP Managers	-	-	X
Auto-MDIX Configuration	X	-	-
BootP	X	-	-
Config File	X	-	-
Console Access	X	-	-
Copy Command	X	-	-
Debug	X	-	-
DHCP Configuration	-	X	-
DHCP/Bootp Operation	X	-	-
DHCP Option 82	-	X	-
Diagnostic Tools	X	-	-
Downloading Software	X	-	-
Event Log	X	-	-
Factory Default Settings	X	-	-
File Management	X	-	-

Feature	Management and Configuration	Advanced Traffic Management	Access Security Guide
File Transfers	X	-	-
Friendly Port Names	X		
GVRP	-	X	-
IGMP	-	X	-
Interface Access (Telnet, Console/Serial, Web)	X	-	-
Jumbo Packets	X	-	-
IP Addressing	X	-	-
IP Routing	-	X	-
LACP	X	-	-
Link	X	-	-
LLDP	X	-	-
LLDP-MED	X	-	-
MAC Address Management	X	-	-
MAC Lockdown	-	-	X
MAC Lockout	-	-	X
MAC-based Authentication	-	-	X
Monitoring and Analysis	X	-	-
Multicast Filtering	-	X	-
Multiple Configuration Files	X	-	-
Network Management Applications (LLDP, SNMP)	X	-	-
Passwords	-	-	X
Ping	X	-	-
Port Configuration	X	-	-
Port Security	-	-	X
Port Status	X	-	-
Port Trunking (LACP)	X	-	-

Product Documentation

Feature	Management and Configuration	Advanced Traffic Management	Access Security Guide
Port-Based Access Control	-	-	X
Port-Based Priority (802.1Q)	X	-	-
Power over Ethernet (PoE)	X	-	-
Quality of Service (QoS)	-	X	-
RADIUS ACLs	-	-	X
RADIUS Authentication and Accounting	-	-	X
Routing	-	X	-
Secure Copy	X	-	-
sFlow	X		
SFTP	X	-	-
SNMP	X	-	-
Software Downloads (SCP/SFTP, TFTP, Xmodem)	X	-	-
Source-Port Filters	-	-	X
Spanning Tree (STP, RSTP, MSTP)	-	X	-
SSH (Secure Shell) Encryption	-	-	X
SSL (Secure Socket Layer)	-	-	X
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	-	-
Uni-Directional Link Detection (UDLD)	X	-	-

Feature	Management and Configuration	Advanced Traffic Management	Access Security Guide
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	1-6
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

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 Series 2610
- ProCurve Series 2610-PWR

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

You can also download the software manuals from the ProCurve website, 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 7-1): Provides encrypted paths for remote access to switch management functions.
- **Secure Socket Layer (SSL)** (page 8-1): Provides remote web access to the switch via encrypted authentication paths between the switch and management station clients capable of SSL/TLS operation.

- **Access Control Lists** (page 9-1): Permits or denies in-band management access. This includes preventing the use of certain TCP or UDP applications (such as Telnet, SSH, Web browser, and SNMP) for transactions between specific source and destination IP addresses. Eliminates unwanted IP, TCP, or UDP traffic by filtering packets where they enter or leave the switch on specific interfaces.
- **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).
- **Port-Based and User-Based Access Control (802.1X)** (page 11-1): 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 12-1): 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.
- **Authorized IP Managers** (page 13-1): 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+ ¹	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 ¹	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>

¹ The local Manager/Operator, TACACS+, and RADIUS options (direct connect or modem access) also offer protection for serial port access.

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 ***Series 2610/2610-PWR Switches***”.

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 2610#
```

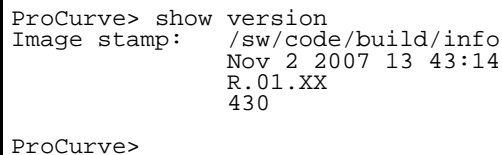
To simplify recognition, this guide uses `ProCurve` to represent command prompts for all models. For example:

```
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> show version
Image stamp:   /sw/code/build/info
               Nov 2 2007 13 43:14
               R.01.XX
               430

ProCurve>
```

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 xiii.

Note

For the latest version of all ProCurve switch documentation, including release notes covering recently added features, visit the ProCurve Networking website at 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:




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:

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 *Installation and Getting Started Guide* you received with the switch.

To Set Up and Install the Switch in Your Network

Important!

Use the *Installation and Getting Started Guide* shipped with your switch for the following:

- Notes, cautions, and warnings related to installing and using the switch and its related modules
- 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 website. (Refer to “Product Documentation” on page xiii of this guide for further details.)

Getting Started
Need Only a Quick Start?

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	2-16
Password Recovery Process	2-18

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.

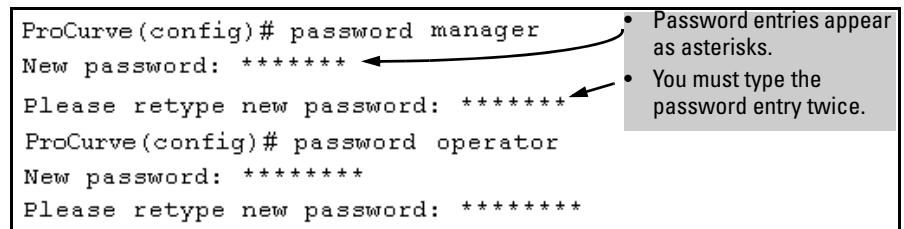
CLI: Setting Passwords and Usernames

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: *****
```

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

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

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

```
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 2-18.) (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	3-15
Configuring Web Authentication	3-18
Overview	3-18
Configure the Switch for Web-Based Authentication	3-19
Configuring MAC Authentication on the Switch	3-23
Overview	3-23
Configure the Switch for MAC-Based Authentication	3-24
Show Commands for Web-Based Authentication	3-28
Show Commands for MAC-Based Authentication	3-31
Show Client Status	3-33

Overview

Feature	Default	Menu	CLI	Web
Configure Web Authentication	n/a	—	3-18	—
Configure MAC Authentication	n/a	—	3-23	—
Display Web Authentication Status and Configuration	n/a	—	3-28	—
Display MAC Authentication Status and Configuration	n/a	—	3-31	—

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

802.1X port-access and either Web authentication or MAC authentication can be concurrently configured on the same port, with a maximum of eight 802.1X clients allowed on the port. (The default is one client.)

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 8 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 include 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.

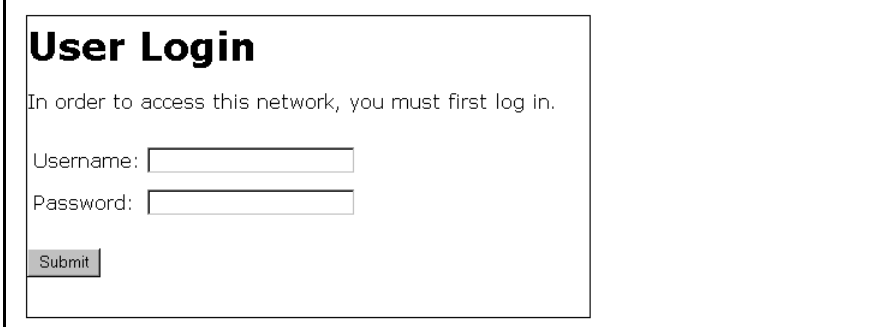
A screenshot of a web-based user login interface. The interface is enclosed in a rectangular border. At the top left, the text "User Login" is displayed in a large, bold, black font. Below this, a message reads "In order to access this network, you must first log in." in a smaller black font. Underneath the message, there are two input fields: the first is labeled "Username:" and the second is labeled "Password:". Both labels are in a small black font. Below the password field, there is a button labeled "Submit" in a small black font. The entire form is set against a white background.

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 their 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.

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 3-22.

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

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

Operating Rules and Notes

- The switch supports concurrent 802.1X and either Web- or MAC-authentication operation on a port (with up to 8 clients allowed). However, concurrent operation of Web- or MAC-authentication with other types of authentication on the same port is not supported. That is, the following authentication types are mutually exclusive on a given port:
 - Web Authentication (with or without 802.1X)
 - MAC Authentication (with or without 802.1X)
 - MAC lockdown
 - MAC lockout
 - Port-Security
- Order of Precedence for Port Access Management (highest to lowest):
 - MAC lockout
 - MAC lockdown or Port Security
 - Port-based Access Control (802.1X) or Web Authentication or MAC Authentication

Note on Port Access Management

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

- VLANs: If your LAN does not use multiple VLANs, then you do not need to configure VLAN assignments in your RADIUS server or consider using either Authorized or Unauthorized VLANs. If your LAN does use multiple VLANs, then some of the following factors may apply to your use of Web-Auth and MAC-Auth.
 - Web-Auth and MAC-Auth operate only with port-based VLANs. Operation with protocol VLANs is not supported, and clients do not have access to protocol VLANs during Web-Auth and MAC-Auth sessions.
 - A port can belong to one, untagged VLAN during any client session. Where multiple authenticated clients may simultaneously use the same port, they must all be capable of operating on the same VLAN.

- During an authenticated client session, the following hierarchy determines a port's VLAN membership:
 1. If there is a RADIUS-assigned VLAN, then, for the duration of the client session, the port belongs to this VLAN and temporarily drops all other VLAN memberships.
 2. If there is no RADIUS-assigned VLAN, then, for the duration of the client session, the port belongs to the Authorized VLAN (if configured) and temporarily drops all other VLAN memberships.
 3. If neither 1 or 2, above, apply, but the port is an untagged member of a statically configured, port-based VLAN, then the port remains in this VLAN.
 4. If neither 1, 2, or 3, above, apply, then the client session does not have access to any statically configured, untagged VLANs and client access is blocked.
 - After an authorized client session begins on a given port, the port's VLAN membership does not change. If other clients on the same port become authenticated with a different VLAN assignment than the first client, the port blocks access to these other clients until the first client session ends.
 - The optional “authorized” VLAN (**auth-vid**) and “unauthorized” VLAN (**unauth-vid**) you can configure for Web- or MAC-based authentication must be statically configured VLANs on the switch. Also, if you configure one or both of these options, any services you want clients in either category to access must be available on those VLANs.
-
- Where a given port's configuration includes an unauthorized client VLAN assignment, the port will allow an unauthenticated client session only while there are no requests for an authenticated client session on that port. In this case, if there is a successful request for authentication from an authorized client, the switch terminates the unauthorized-client session and begins the authorized-client session.
 - When a port on the switch is configured for Web or MAC Authentication and is supporting a current session with another device, rebooting the switch invokes a re-authentication of the connection.
 - When a port on the switch is configured as a Web- or MAC-based authenticator, it blocks access to a client that does not provide the proper authentication credentials. If the port configuration includes an optional, unauthorized VLAN (**unauth-vid**), the port is temporarily placed in the unauthorized VLAN if there are no other authorized clients currently using the port with a different VLAN assignment. If an authorized client is using the port with a different VLAN or if there is no unauthorized VLAN configured, the unauthorized client does not receive access to the network.

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

**Note on Web/
MAC
Authentication
and LACP**

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

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 eight format options:

aabbccddeeff (the default format)

aabbcc-ddeeff

aa-bb-cc-dd-ee-ff

aa:bb:cc:dd:ee:ff

AABBCCDDEEFF

AABBCC-DDEEFF

AA-BB-CC-DD-EE-FF

AA:BB:CC:DD:EE:FF

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

Configuring the Switch To Access a RADIUS Server

RADIUS Server Configuration Commands

radius-server	
[host <ip-address> [auth-port UDP-PORT acct-port UDP- PORT]]	below
[key < global-key-string >]	below
timeout	3-16
retransmit	3-16
dead-time	3-16
radius-server host <ip-address> key <server-specific key-string>	3-16

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.)

timeout <1-15>

*The server response timeout interval in seconds.
Default: 5 seconds*

retransmit <1-5>

Specifies the maximum number of retransmission attempts. Default: 3 attempts

dead-time <1-1440> (in minutes)

If the switch does not receive a response from a specific RADIUS server, the switch does not send any new authentication requests to that server until the dead-time has expired. During a new authentication attempt, the switch bypasses a specified RADIUS server if a dead-time period is running on the switch because of a previous failure to receive a response from that server. The switch continues to send new authentication requests to any other configured RADIUS servers not affected by a dead-time condition.

Dead-time begins with the end of the last timeout in the last retransmit attempt of the failed authentication session. When dead-time is set to zero, there is no dead-time and the switch will not bypass a RADIUS server that has failed to respond to an earlier authentication attempt.

Default: 0 (zero)

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

Optional. Specifies an encryption key for use during authentication (or accounting) sessions with the 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 Port	Acct Port	Encryption Key
-----	----	----	-----
192.168.32.11	1812	1813	2Pzo22

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

no aaa port-access web-based < 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 < port-list> [client-limit <1-8>]

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

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

Forces a reauthentication of all attached clients on the port.

Syntax: aaa port-access web-based < port-list> [redirect-url <url>]
no aaa port-access web-based < 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 < 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 < 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 < port-list> [unauth-vid <vid>]
no aaa port-access web-based < port-list> [unauth-vid]

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

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

Configuring MAC Authentication on the Switch

Overview

1. If you have not already done so, configure a local username and password pair on the switch.
2. If you plan to use multiple VLANs with MAC Authentication, ensure that these VLANs are configured on the switch and that the appropriate port assignments have been made.
3. Use the **ping** command in the switch console interface to ensure that the switch can communicate with the RADIUS server you have configured to support MAC-Auth on the switch.
4. Configure the switch with the correct IP address and encryption key to access the RADIUS server.
5. Configure the switch for MAC-Auth:
 - a. Configure MAC Authentication on the switch ports you want to use.
6. Test both the authorized and unauthorized access to your system to ensure that MAC Authentication works properly on the ports you have configured for port-access.

Configure the Switch for MAC-Based Authentication

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

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

Specifies the MAC address format to be used in the RADIUS request message. This format must match the format used to store the MAC addresses in the RADIUS server. (Default: no-delimiter)

no-delimiter — specifies an aabbccddeeff format.

single-dash — specifies an aabbcc-ddeeff format.

multi-dash — specifies an aa-bb-cc-dd-ee-ff format.

multi-colon — specifies an aa:bb:cc:dd:ee:ff format.

no-delimiter-uppercase—specifies an AABBCCDDEEFF format

single-dash-uppercase—specifies an AABBCD-DDEEFF format

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

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

Syntax: [no] aaa port-access mac-based < 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 < port-list > [addr-limit <1-8>]

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

Syntax: [no] aaa port-access mac-based < 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 < port-list > [auth-vid <vid>]

no aaa port-access mac-based < 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 < *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 < *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 < *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 < *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 < *port-list* > [reauthenticate]

Forces a reauthentication of all attached clients on the port.

Syntax: aaa port-access mac-based < 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 < *port-list* > [unauth-vid < *vid* >]
 no aaa port-access mac-based < *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 Commands for Web-Based Authentication

Command	Page
show port-access [<i>port-list</i>] web-based	3-28
[clients]	3-28
[config]	3-28
[config [auth-server]]	3-29
[config [web-server]]	3-29
show port-access <i>port-list</i> web-based config detail	3-29

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.

Example: Verifying a Web Authentication Configuration

The following example shows how to use the **show port-access web-based config** command to display the currently configured web-authentication settings for all switch ports, including:

- Temporary DHCP base address and mask
- Authorized and unauthorized VLAN IDs
- Controlled directions setting for transmitting Wake-on-LAN traffic on egress ports

Web and MAC Authentication

Show Commands for Web-Based Authentication

```
ProCurve(config)# show port-access web-based config

Port Access Web-Based Configuration

DHCP Base Address : 192.168.0.0
DHCP Subnet Mask  : 255.255.255.0
DHCP Lease Length : 10

Allow RADIUS-assigned dynamic (GVRP) VLANs [No] : No
```

Port	Enabled	Client Limit	Client Moves	Logoff Period	Re-Auth Period	Unauth VLAN ID	Auth VLAN ID	Cntrl Dir
1	No	1	No	300	0	0	0	both
2	No	1	No	300	0	0	0	both
3	No	1	No	300	0	0	0	both
4	No	1	No	300	0	0	0	both
5	No	1	No	300	0	0	0	both
6	No	1	No	300	0	0	0	both
7	No	1	No	300	0	0	0	both
8	No	1	No	300	0	0	0	both

Figure 3-5. Example of Verifying a Web Authentication Configuration

Show Commands for MAC-Based Authentication

Command	Page
show port-access [<i>port-list</i>] mac-based	3-31
[clients]	3-31
[config]	3-31
[config [auth-server]]	3-32
show port-access <i>port-list</i> mac-based config detail	3-32

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

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

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

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

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

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

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

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

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

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

Example: Verifying a MAC Authentication Configuration

The following example shows how to use the **show port-access mac-based config** command display the currently configured MAC authentication settings for all switch ports, including:

- MAC address format
- Authorized and unauthorized VLAN IDs
- Controlled directions setting for transmitting Wake-on-LAN traffic on egress ports

```
ProCurve(config)# show port-access mac-based config

Port Access MAC-Based Configuration

MAC Address Format : no-delimiter
Allow RADIUS-assigned dynamic (GVRP) VLANs [No] : No

  Port  Enabled  Client Limit  Client Moves  Logoff Period  Re-Auth Period  Unauth VLAN ID  Auth VLAN ID  Cntrl Dir
  ----  -
  1      No       1           No            300           0               0               0             both
  2      No       1           No            300           0               0               0             both
  3      No       1           No            300           0               0               0             both
  4      No       1           No            300           0               0               0             both
  5      No       1           No            300           0               0               0             both
  6      No       1           No            300           0               0               0             both
```

Figure 3-6. Example of Verifying a MAC Authentication Configuration

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.

Web and MAC Authentication

Show Client Status

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	4-10
Configuring the Switch's TACACS+ Server Access	4-17
How Authentication Operates	4-22
General Authentication Process Using a TACACS+ Server	4-22
Local Authentication Process	4-24
Using the Encryption Key	4-25
Controlling Web Browser Interface Access When Using TACACS+ Authentication	4-26
Messages Related to TACACS+ Operation	4-27
Operating Notes	4-28

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 4-10	—
configure the switch to contact TACACS+ server(s)	disabled	—	page 4-17	—

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).

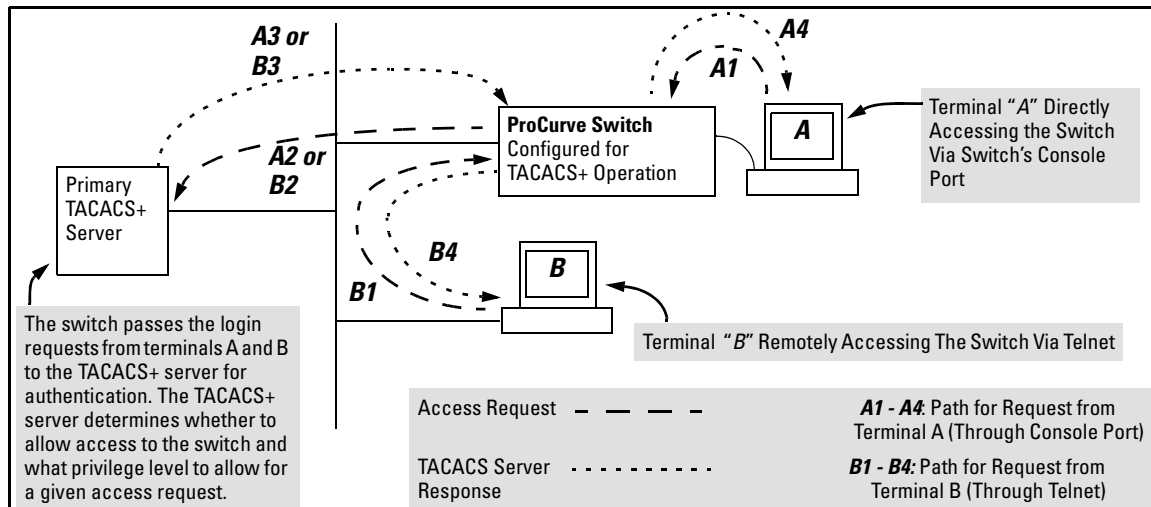


Figure 4-1. Example of TACACS+ Operation

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 authentication.

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 4-26.

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 4-26.

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 4-25.)
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 4-10 through 4-16
console	
Telnet	
num-attempts <1-10 >	
login <privilege-mode>	
tacacs-server	pages 4-17
host < ip-addr >	pages 4-17
key	4-21
timeout < 1-255 >	4-22

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					
Access Task	Login Primary	Login Secondary	Enable Primary	Enable Secondary	
-----	-----	-----	-----	-----	
(Console _ _ _	Local _ _ _	None _ _ _	Local _ _ _	None _ _ _	Configuration for login and enable access to the switch through the switch console port.
(Telnet _ _ _	Local _ _ _	None _ _ _	Local _ _ _	None _ _ _	

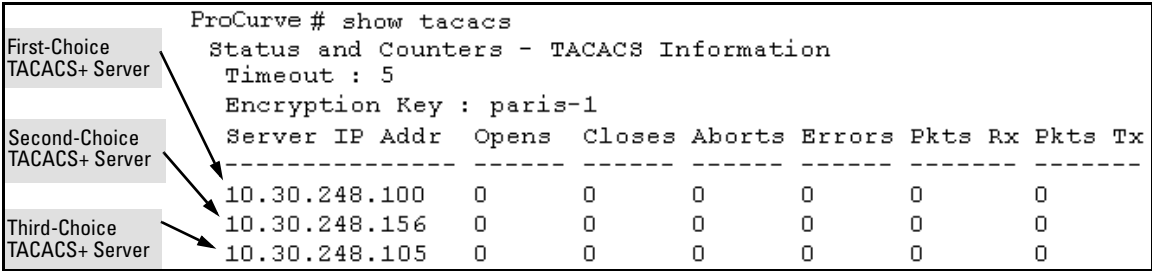
Figure 4-2. Example Listing of the Switch's Authentication Configuration

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
```

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

Configuring the Switch's Authentication Methods

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

- Console
- Telnet
- SSH
- Web
- Port-access (802.1X)

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

Using the Privilege-Mode Option for Login

When using TACACS+ to control user access to the switch, you must first login with your username at the Operator privilege level using the password for Operator privileges, and then login again with the same username but using the Manager password to obtain Manager privileges. You can avoid this double login process by entering the **privilege-mode** option with the **aaa authentication login** command to enable TACACS+ for a single login. The switch authenticates your username/password, then requests the privilege level (Operator or Manager) that was configured on the TACACS+ server for this username/password. The TACACS+ server returns the allowed privilege level to the switch. You are placed directly into Operator or Manager mode, depending on your privilege level.

```
ProCurve(config) aaa authentication login privilege-mode
```

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

Syntax: aaa authentication

< console | telnet | ssh >

Selects the access method for configuration.

< enable >

The server grants privileges at the Manager privilege level.

< login [privilege-mode] >

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

Default: Single login disabled.

< local | tacacs | radius >

Selects the type of security access:

local — Authenticates with the Manager and Operator password you configure in the switch.

tacacs — Authenticates with a password and other data configured on a TACACS+ server.

radius — Authenticates with a password and other data configured on a RADIUS server.

[< local | none >]

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

aaa authentication num-attempts < 1-10 >

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

Table 4-1. AAA Authentication Parameters

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

Configuring the TACACS+ Server for Single Login

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

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

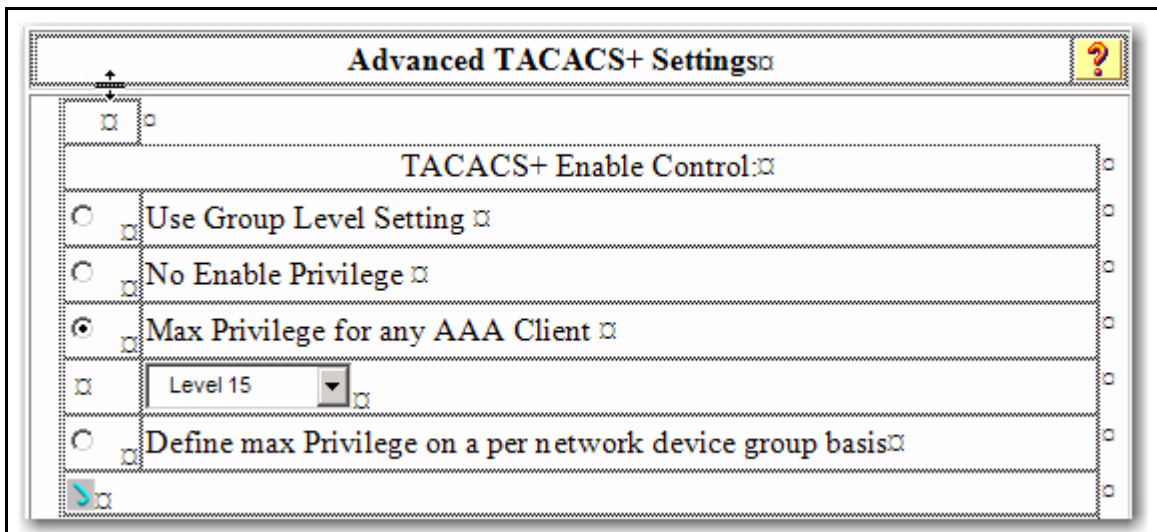


Figure 4-4. Advanced TACACS+ Settings Section of the TACACS+ Server User Setup

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

Check the Privilege level box and set the privilege level to 15 to allow “root” privileges. This allows you to use the single login option.

<input checked="" type="checkbox"/>	Shell (exec) <input type="button" value="X"/>	
<input type="checkbox"/>	Access control list <input type="button" value="X"/>	<input type="text" value=""/>
<input type="checkbox"/>	Auto command <input type="button" value="X"/>	<input type="text" value=""/>
<input type="checkbox"/>	Callback line <input type="button" value="X"/>	<input type="text" value=""/>
<input type="checkbox"/>	Callback rotary <input type="button" value="X"/>	<input type="text" value=""/>
<input type="checkbox"/>	Idle time <input type="button" value="X"/>	<input type="text" value=""/>
<input type="checkbox"/>	No callback verify <input type="button" value="X"/>	<input type="checkbox"/> Enabled <input type="button" value="X"/>
<input type="checkbox"/>	No escape <input type="button" value="X"/>	<input type="checkbox"/> Enabled <input type="button" value="X"/>
<input type="checkbox"/>	No hangup <input type="button" value="X"/>	<input type="checkbox"/> Enabled <input type="button" value="X"/>
<input checked="" type="checkbox"/>	Privilege level <input type="button" value="X"/>	<input type="text" value="15"/>
<input type="checkbox"/>	Timeout <input type="button" value="X"/>	<input type="text" value="123"/>
<input type="checkbox"/>	Custom attributes <input type="button" value="X"/>	

Figure 4-5. The Shell Section of the TACACS+ Server User Setup

Primary/Secondary Authentication

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.

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 Addr list, if any. If the switch still fails to receive a response from any TACACS+ server, it reverts to whatever secondary authentication method was configured using the **aaa authentication** command (local or none; see “Configuring the Switch's Authentication Methods” on page 4-10.)

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 4-25 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 4-22.

Name	Default	Range
[key <key-string>]	none (null)	n/a
Specifies the optional, global “encryption key” that is also assigned in the TACACS+ server(s) that the switch will access for authentication. This option is subordinate to any “per-server” encryption keys you assign, and applies only to accessing TACACS+ servers for which you have not given the switch a “per-server” key. (See the host <ip-addr> [key <key-string>] entry at the beginning of this table.) For more on the encryption key, see “Using the Encryption Key” on page 4-25 and the documentation provided with your TACACS+ server application.		
timeout <1 - 255>	5 sec	1 - 255 sec
Specifies how long the switch waits for a TACACS+ server to respond to an authentication request. If the switch does not detect a response within the timeout period, it initiates a new request to the next TACACS+ server in the list. If all TACACS+ servers in the list fail to respond within the timeout period, the switch uses either local authentication (if configured) or denies access (if none configured for local authentication).		

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

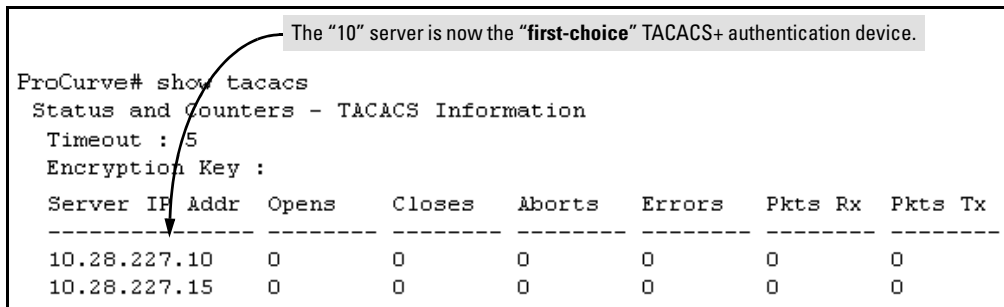
Suppose that the switch was already configured to use TACACS+ servers at 10.28.227.10 and 10.28.227.15. In this case, 10.28.227.15 was entered first, and so is listed as the first-choice server:

ProCurve# show tacacs								
Status and Counters - TACACS Information								
Timeout : 5								
Encryption Key :								
First-Choice TACACS+ Server								
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 4-6. Example of the Switch with Two TACACS+ Server Addresses Configured

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

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



```
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 4-7. Example of the Switch After Assigning a Different “First-Choice” Server

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

```
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 4-25.)

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.

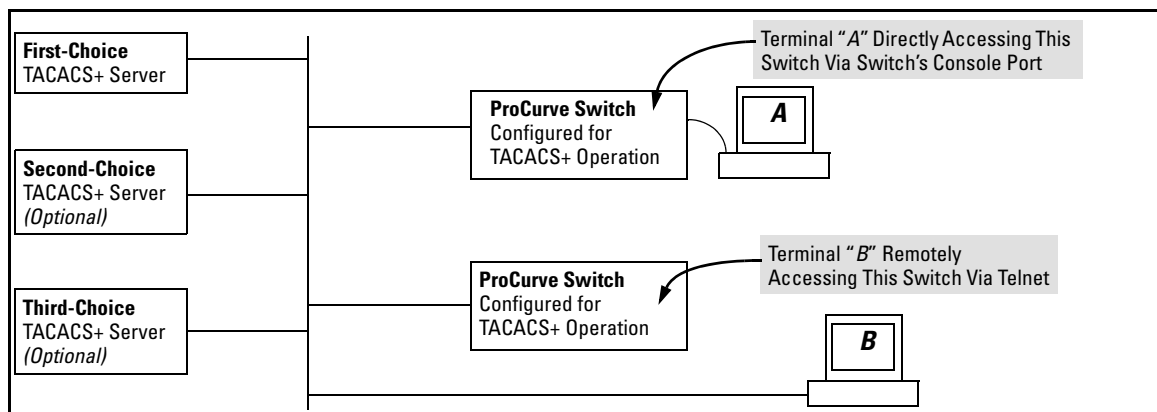


Figure 4-8. Using a TACACS+ Server for Authentication

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

1. The switch queries the first-choice TACACS+ server for authentication of the request.
 - If the switch does not receive a response from the first-choice TACACS+ server, it attempts to query a secondary server. If the switch does not receive a response from any TACACS+ server, then it uses its own local username/password pairs to authenticate the logon request. (See "Local Authentication Process" on page 4-24.)
 - 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 4-15.)

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	5-11
3. Configure the Switch's Global RADIUS Parameters	5-13
Local Authentication Process	5-17
Controlling Web Browser Interface Access When Using RADIUS Authentication	5-18
Configuring RADIUS Authorization	5-18
Overview	5-18
Commands Authorization Type	5-19
Enabling Authorization with the CLI	5-19
Showing Authorization Information	5-20
Configuring the RADIUS Server	5-20
Configuring RADIUS Accounting	5-26
Operating Rules for RADIUS Accounting	5-27
Steps for Configuring RADIUS Accounting	5-28
Viewing RADIUS Statistics	5-33
General RADIUS Statistics	5-33
RADIUS Authentication Statistics	5-35
RADIUS Accounting Statistics	5-36

Changing RADIUS-Server Access Order	5-37
Messages Related to RADIUS Operation	5-39

Overview

Feature	Default	Menu	CLI	Web
Configuring RADIUS Authentication	None	n/a	5-6	n/a
Configuring RADIUS Accounting	None	n/a	5-26	n/a
Viewing RADIUS Statistics	n/a	n/a	5-33	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

For information on blocking unauthorized access through the web browser interface, refer to “Controlling Web Browser Interface Access When Using RADIUS Authentication” on page 5-18.

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 5-33). 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 5-37.)
- 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 authorized> <web-based mac-based> <chap-radius>>>	5-8
[no] radius-server host < IP-address >	5-11
[auth-port < port-number >]	5-11
[acct-port < port-number >]	5-11, 5-28
[key < server-specific key-string >]	5-11
[no] radius-server key < global key-string >	5-13
radius-server timeout < 1 - 15>	5-13
radius-server retransmit < 1 - 5 >	5-13
[no] radius-server dead-time < 1 - 1440 >	5-15
show radius	5-33
[< host < ip-address>]	5-33
show authentication	5-35
show radius authentication	5-35

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 5-26.)

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 7-1.
- **Web:** Web browser interface.

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

Syntax: `aaa authentication < console | telnet | ssh | web > < enable | login > < local | radius>>>< web-based | mac-based> <chap-radius>>`

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

`<console | telnet | ssh | web>`

`[< local | none | authorized >]`

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

*The **authorized** option allows users unconditional access to the network when the primary authentication method fails. See **Caution** below.*

`<web-based | mac-based> <chap-radius>>`

Password authentication for web-based or mac-based port access to the switch.

`[none | authorized]`

*Provides options for secondary authentication. The **none** option specifies that a backup authentication method is not used. The **authorized** option allows access without authentication.*

*Default: **none***

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

Caution

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

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):

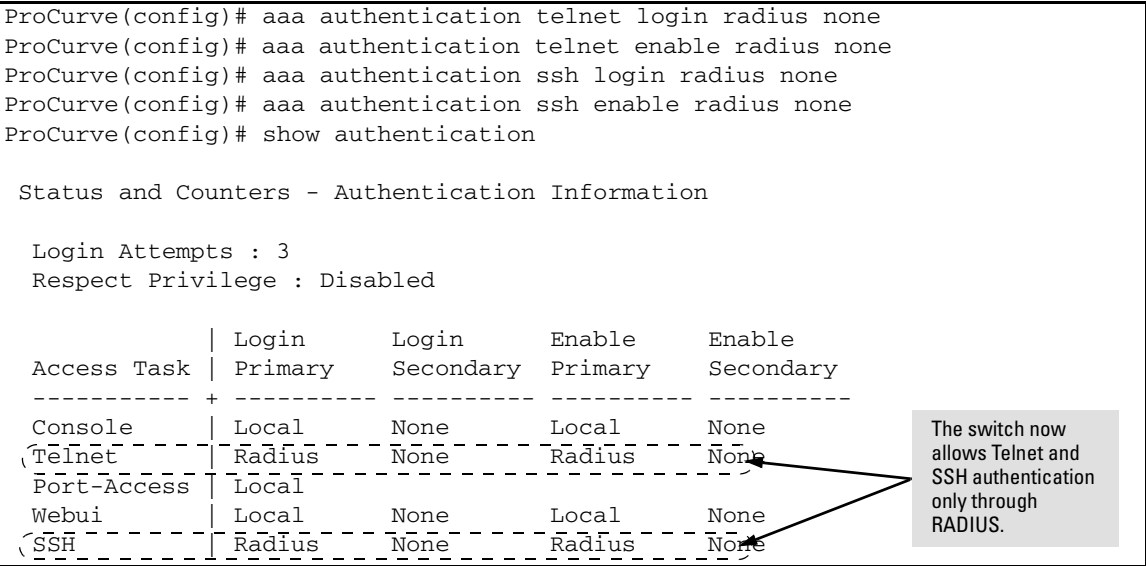


Figure 5-2. Example Configuration for RADIUS Authentication

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 5-17.

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 5-26: “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 5-37.)*

[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 5-37.

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 5-11.)
- **Server timeout:** Defines the time period in seconds for authentication attempts. If the timeout period expires before a response is received, the attempt fails.
- **Server dead time:** Specifies the time in minutes during which the switch avoids requesting authentication from a server that has not responded to previous requests.
- **Retransmit attempts:** If the first attempt to contact a RADIUS server fails, specifies how many retries you want the switch to attempt on that server.

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 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
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 RADIUS authentication access.
 - 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 Authorization

Overview

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

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

Commands Authorization Type

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

Note

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

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

Enabling Authorization with the CLI

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

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

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

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

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

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

```
ProCurve(config)# aaa authorization commands radius
```

When the NAS sends the RADIUS server a valid username and password, the RADIUS server sends an Access-Accept packet that contains two attributes—the command list and the command exception flag. When an authenticated user enters a command on the switch, the switch examines the list of commands delivered in the RADIUS Access-Accept packet as well as the command exception flag, which indicates whether the user has permission to execute the commands in the list. See *Configuring the RADIUS Server* on page 5-20.

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

Showing Authorization Information

You can show the authorization information by entering this command:

Syntax: show authorization

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

An example of the output is shown.

```
ProCurve(config)# show authorization

Status and Counters - Authorization Information

Type      | Method
-----+-----
Commands  | RADIUS
```

Figure 5-7. Example of Show Authorization Command

Configuring the RADIUS Server

Using Vendor Specific Attributes (VSAs)

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

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

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

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

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

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

Example Configuration on Cisco Secure ACS for MS Windows

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

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

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

```
[User Defined Vendor]
;
; The Name and IETF vendor code and any VSAs MUST be unique.
;
; One or more VSAs named (max 255)
;
; Each named VSA requires a definition section...
;
; Types are STRING, INTEGER, IPADDR
;
; The profile specifies usage, IN for accounting, OUT for
  authorization,
; MULTI if more than a single instance is allowed per
  RADIUS message.
; Combinations are allowed, e.g. "IN", "MULTI OUT",
  "MULT IN OUT"
;
; Enumerations are optional for INTEGER attribute types

[User Defined Vendor]

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

[Hp-Command-String]
```

```
Type=STRING
Profile=IN OUT

[Hp-Command-Exception]

Type=INTEGER
Profile=IN OUT

Enums=Hp-Command-Exception-Types

[Hp-Command-Exception-Types]

0=PermitList
1=DenyList
```

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

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

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

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

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

```
Parsing [.\hp.ini] for addition at UDV slot [0]
```

```
Stopping any running services
```

```
Creating backup of current config
```

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

```
Done
```

```
Checking new configuration...
```

```
New configuration OK
```

```
Re-starting stopped services
```

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

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

5. Go to:

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

6. Right click and then select **New > key**. Add the vendor Id number that you determined in step 4 (100 in the example).
7. Restart all Cisco services.
8. The newly created HP RADIUS VSA appears only when you configure an AAA client (NAS) to use the HP VSA RADIUS attributes. Select Network Configuration and add (or modify) an AAA entry. In the Authenticate Using field choose RADIUS(HP) as an option for the type of security control protocol.
9. Select **Submit + Restart** to effect the change. The HP RADIUS VSA attributes will appear in Cisco ACS configurations, for example, "Interface Configuration", "Group Setup", "User Setup".

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

1. Select **System Configuration**.
2. Select **Logging**.
3. Select **CSV RADIUS Accounting**. In the Select Columns to Log section, add the HP-Command-String attribute to the Logged Attributes list.
4. Select **Submit**.
5. Select **Network Configuration**. In the AAA Clients section, select an entry in the AAA Client Hostname column. You will go to the AAA Client Setup screen.
6. Check the box for **Log Update/Watchdog Packets from this AAA Client**.
7. Click **Submit + Restart**. You should be able to see the HP-Command-String attribute in the RADIUS accounting reports.

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

Example Configuration Using FreeRADIUS

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

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

VENDOR            Hp            11

# HP Extensions

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

# Hp-Command-Exception Attribute Values

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

2. Find the location of the dictionary files used by FreeRADIUS (try /usr/local/share/freeradius).
3. Copy dictionary.hp to that location. Open the existing dictionary file and add this entry:
\$ INCLUDE dictionary.hp
4. You can now use HP VSAs with other attributes when configuring user entries.

Configuring RADIUS Accounting

RADIUS Accounting Commands	Page
[no] radius-server host < ip-address >	5-28
[acct-port < port-number >]	5-28
[key < key-string >]	5-28
[no] aaa accounting < exec network system >	5-31
< start-stop stop-only> radius	
[no] aaa accounting update	5-32
periodic < 1 - 525600 > (in minutes)	
[no] aaa accounting suppress null-username	5-32
show accounting	5-36
show accounting sessions	5-37
show radius accounting	5-36

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

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

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

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

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

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

Operating Rules for RADIUS Accounting

- You can configure up to three types of accounting to run simultaneously: exec, system, and network.
- RADIUS servers used for accounting are also used for authentication.
- The switch must be configured to access at least one RADIUS server.
- RADIUS servers are accessed in the order in which their IP addresses were configured in the switch. Use **show radius** to view the order. As long as the first server is accessible and responding to authentication requests from the switch, a second or third server will not be accessed. (For more on this topic, refer to “Changing RADIUS-Server Access Order” on page 5-37.)
- 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 5-11.
- Provide the following:
 - A RADIUS server IP address.
 - Optional—a UDP destination port for authentication requests. Otherwise the switch assigns the default UDP port (1812; recommended).
 - Optional—if you are also configuring the switch for RADIUS authentication, and need a unique encryption key for use during authentication sessions with the RADIUS server you are designating, configure a server-specific key. This key overrides the global encryption key you can also configure on the switch, and must match the encryption key used on the specified RADIUS server. For more information, refer to the **key <key-string>** parameter on page 5-11. (Default: null)

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

- **Accounting types:** exec (page 5-27), network (page 5-26), or system (page 5-27)
- **Trigger for sending accounting reports to a RADIUS server:** At session start and stop or only at session stop

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

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

1. Configure the Switch To Access a RADIUS Server

Before you configure the actual accounting parameters, you should first configure the switch to use a RADIUS server. This is the same as the process described on page 5-11. You need to repeat this step here only if you have not yet configured the switch to use a RADIUS server, your server data has

changed, or you need to specify a non-default UDP destination port for accounting requests. Note that switch operation expects a RADIUS server to accommodate both authentication and accounting.

Syntax: [no] radius-server host < ip-address >

*Adds a server to the RADIUS configuration or (with **no**) deletes a server from the configuration.*

[acct-port < port-number >]

Optional. Changes the UDP destination port for accounting requests to the specified RADIUS server. If you do not use this option, the switch automatically assigns the default accounting port number. (Default: 1813)

[key < key-string >]

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

(For a more complete description of the **radius-server** command and its options, turn to page 5-11.)

For example, suppose you want 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.).

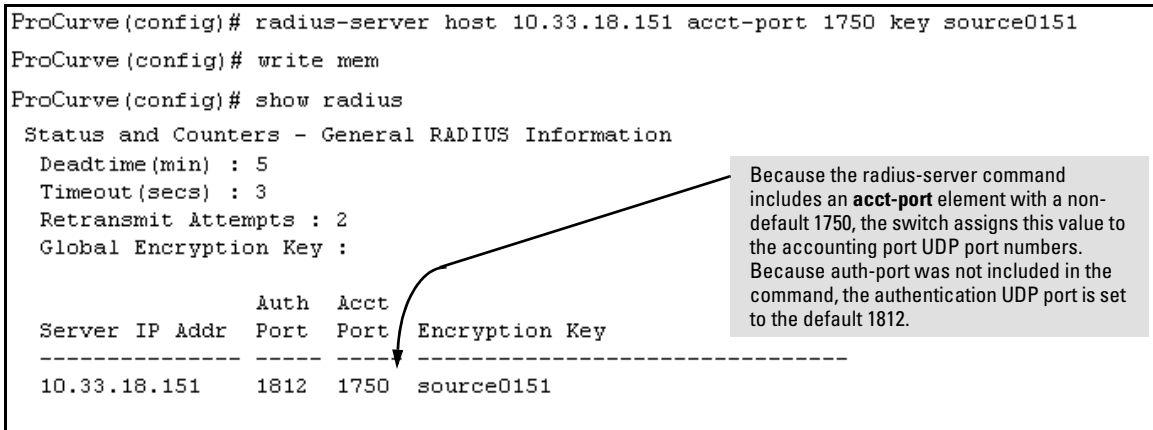


Figure 5-8. Example of Configuring for a RADIUS Server with a Non-Default Accounting UDP Port Number

The radius-server command as shown in figure 5-8, 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.)

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 5-30) 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 acknowledgement.

The system option (page 5-30) 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

Figure 5-9 includes three callout boxes with arrows pointing to specific parts of the terminal output:

- Configures exec and system accounting and controls.** (Points to the two configuration commands)
- Summarizes the switch's accounting configuration.** (Points to the 'show accounting' command)
- Exec and System accounting are active. (Assumes the switch is configured to access a reachable** (Points to the 'Exec' and 'System' rows in the table)

Figure 5-9. 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 5-9, 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 ← Update Period

Suppress Empty User : No ← Suppress Unknown User

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

Figure 5-10. 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 5-26.)*

```
ProCurve(config)# show radius
Status and Counters - General RADIUS Information
  Deadttime(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 5-11. 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
  Round Trip Time : 2
  Pending Requests : 0
  Retransmissions : 0
  Timeouts : 0
  Malformed Responses : 0
  Bad Authenticators : 0
  Unknown Types : 0
  Packets Dropped : 0
  Access Requests : 2
  Access Challenges : 0
  Access Accepts : 2
  Access Rejects : 0
  Accounting UDP Port : 1813
  Round Trip Time : 7
  Pending Requests : 0
  Retransmissions : 0
  Timeouts : 0
  Malformed Responses : 0
  Bad Authenticators : 0
  Unknown Types : 0
  Packets Dropped : 0
  Accounting Requests : 2
  Accounting Responses : 2
```

Figure 5-12. RADIUS Server Information From the Show Radius Host Command

Table 5-2. Values for Show Radius Host Output (Figure 5-12)

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

RADIUS Authentication Statistics

Syntax: show authentication

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

show radius authentication

*Displays NAS identifier and data on the configured RADIUS server and the switch's interactions with this server. (Requires prior use of the **radius-server host** command to configure a RADIUS server IP address in the switch. See "Configuring RADIUS Accounting" on page 5-26.)*

```
ProCurve> show authentication
Status and Counters - Authentication Information
Login Attempts : 2
```

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

Figure 5-13. 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
```

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

Figure 5-14. 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 5-15. 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 5-16. 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 5-17. 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  Acct
                Port  Port  Encryption Key
-----
10.10.10.1      1812 1813
10.10.10.2      1812 1813
10.10.10.3      1812 1813
```

RADIUS server IP addresses listed in the order in which the switch will try to access them. In this case, the server at IP address 10.10.10.1 is first.

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

Figure 5-18. 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
Deadtime(min) : 0
Timeout(secs) : 5
Retransmit Attempts : 3
Global Encryption Key :
```

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.

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	

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

Figure 5-19. 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 RADIUS Server Support for Switch Services

Contents

Overview	6-2
Configuring the RADIUS Server for CoS Services	6-3
Viewing the Currently Active Per-Port CoS Configuration Specified by a RADIUS Server	6-3
Configuring and Using RADIUS-Assigned Access Control Lists	6-6
Introduction	6-6
Terminology	6-6
Overview of RADIUS-Assigned, Dynamic Port ACLs	6-9
Contrasting Dynamic and Static ACLs	6-11
How a RADIUS Server Applies a Dynamic Port ACL to a Switch Port	6-12
General ACL Features, Planning, and Configuration	6-13
The Packet-filtering Process	6-14
Operating Rules for Dynamic Port ACLs	6-14
Configuring an ACL in a RADIUS Server	6-15
Configuring ACE Syntax in RADIUS Servers	6-18
Configuring the Switch To Support Dynamic Port ACLs	6-20
Displaying the Current Dynamic Port ACL Activity on the Switch	6-21
Event Log Messages	6-24
Causes of Client Deauthentication Immediately After Authenticating	6-25
Monitoring Shared Resources	6-25

Overview

This chapter provides information that applies to setting up a RADIUS server to configure the following switch features on ports supporting RADIUS-authenticated clients:

- CoS
- ACLS

Optional Network Management Applications. CoS assignments through a RADIUS server are also supported in the ProCurve Manager (PCM) application. ACLs through a RADIUS server can also be augmented using the Identity-Driven Management (IDM) application available for use with PCM. However, the features described in this chapter can be used without PCM or IDM support, if desired.

For information on configuring client authentication on the switch, refer to Chapter 5, “Configuring RADIUS Server Support for Switch Services”.

Optional PCM and IDM Applications. ProCurve Manager is a Windows-based network management solution for all manageable ProCurve devices. It provides network: mapping and polling capabilities, device auto-discovery and topology, tools for device configuration and management, monitoring network traffic, and alerts and troubleshooting information for ProCurve networks.

ProCurve Identity Driven Manager (IDM) is an add-on module to the ProCurve Manager plus (PCM+) application. IDM extends the functionality of PCM+ to include authorization control features for edge devices in networks using RADIUS servers and Web-Authentication, MAC-Authentication, or 802.1X security protocols.

For more information, including electronic copies of the PCM and IDM manuals, visit the ProCurve Web site at www.procurve.com. (The PCM and IDM documentation is available under **Network Management** on the **Product manuals page** of the **Technical Support** area.)

Configuring the RADIUS Server for CoS Services

This section provides general guidelines for configuring a RADIUS server to dynamically apply CoS (Class of Service) for inbound traffic on ports supporting authenticated clients. To configure support for these services on a specific RADIUS server application, refer to the documentation provided with the application. (If multiple clients are authenticated on a port where inbound CoS values have been imposed by a RADIUS server, the CoS applied to all clients on the port are those that are assigned by RADIUS for the most recently authenticated client. Refer to the Note on page 6-5.)

Service	Control Method and Operating Notes:
802.1p (CoS) Priority Assignments on Inbound Traffic This feature assigns a RADIUS-specified 802.1p priority to all inbound packets received on a port supporting an authenticated client.	Vendor-Specific Attribute configured in the RADIUS server. ProCurve (HP) vendor-specific ID:11 VSA: 40 (string = HP) Setting: HP-COS = xxxxxxxx where: x = desired 802.1p priority Note: This is typically an eight-octet field. Enter the same x-value in all eight octets Requires a port-access (802.1X Web Auth, or MAC Auth) authentication method configured on the client's port on the ProCurve switch. For more on 802.1p priority levels, refer to the section titled "Overview" in the "Quality of Service (QoS)" chapter of the <i>Advanced Traffic Management Guide</i> for your switch.

Viewing the Currently Active Per-Port CoS Configuration Specified by a RADIUS Server

While a port-access authenticated client session is active, any RADIUS-imposed port settings override their counterparts in the port's configuration.

Syntax: show port-access authenticator [port-list]
show qos port-priority

These commands display the CoS settings specified by the RADIUS server used to grant authentication for a given client on a given port. When the authenticated client session closes, the switch resets these fields to the values to which they are configured in the switch's running-config file.

show port-access authenticator [port-list] displays, for 802.1X authentication, the status of RADIUS-imposed overrides of the switch's per-port CoS configuration.

show qos port-priority displays, for all port-access authentication methods (802.1X, Web-Auth, and MAC-Auth), the status of RADIUS-imposed overrides of the switch's per-port CoS (802.1p) priority for inbound packets.

ProCurve(config)# show port-access authenticator

Port Access Authenticator Status

Port-access authenticator activated [No] : Yes

Port	Status	Current VLAN ID	Current Port COS	RADIUS ACL Applied?
B7	Open	1	No-override	No
B8	Closed	1	No-override	No
B9	Open	1	7	Yes
B10	Closed	1	No-override	No

Open indicates that there is an authenticated client session running on port B7. **No-override** indicates that there are no RADIUS-imposed settings for CoS (802.1p priority) on port B7.

Open indicates that there is an authenticated client session running on port B9. The numeric value **7** is the most recent RADIUS-imposed setting for the CoS (802.1p priority) on port B9.

Figure 6-1. Example of Displaying Inbound CoS Settings Imposed by a RADIUS Session

```
ProCurve(config)# show qos port-priority
```

Port priorities				
Port	Apply rule	DSCP	Priority	Radius Override
B1	Priority		3	No-override
B2	No-override		No-override	No-override
B3	No-override		No-override	No-override
B4	DSCP	001010	2	5
B5	No-override		No-override	No-override
B6	No-override		No-override	No-override

Priority in the **Apply Rule** column indicates a non-default CoS (802.1p) priority configured in the switch for port B1. The **3** in the **Priority** column shows the actual value configured. **No-override** indicates that there is currently no RADIUS-imposed CoS priority affecting the port.

The **DSCP** in the **Apply Rule** column and the **001010** in the **DSCP** column indicate a non-default CoS (802.1p) priority configured in the switch for packets with a Diffserv codepoint of 001010 inbound on port B4. The **2** in the **Priority** column shows the CoS priority most recently configured for application to packets with that codepoint. The **5** in the **Radius Override** column indicates that there is currently at least one authenticated-client session on port B4, and that the most recent RADIUS-imposed CoS priority for the port is 5, which overrides the configured DSCP setting. See also the following **Note**.

Figure 6-2. Example of Displaying Inbound CoS (802.1p) Priority Imposed by a RADIUS Session

Note

Where multiple clients are currently authenticated on a given port where inbound CoS values have been imposed by a RADIUS server, the port operates with the inbound CoS priority assigned by RADIUS for the most recently authenticated client. Any earlier CoS values on the same port for authenticated client sessions that are still active are overwritten by the most recent RADIUS-imposed values. For example, if client “X” is authenticated with a CoS of 5 and client “Y” later becomes authenticated with a CoS of 3 while the session for client “X” is still active, then the port will operate with a CoS of 3 for both clients.

Configuring and Using RADIUS-Assigned Access Control Lists

Introduction

A RADIUS-assigned ACL is a *dynamic port ACL* configured on a RADIUS server and assigned by the server to filter traffic entering the switch through a specific port from an authenticated client. Note that client authentication can be enhanced by using ProCurve Manager with the optional IDM application. (Refer to “Optional PCM and IDM Applications” on page 6-2.)

The information in this section describes how to apply RADIUS-assigned, dynamic port ACLs on the switch, and assumes a general understanding of ACL structure and operation. If you need information on ACL filtering criteria, design, and operation, please refer to chapter 9, “Access Control Lists (ACLs)”.

Terminology

ACE: See Access Control Entry, below.

Access Control Entry (ACE): An ACE is a policy consisting of a packet-handling action and criteria to define the packets on which to apply the action. For dynamic port ACLs, the elements composing the ACE include:

- **permit** or **drop** (action)
- **in** < *ip-packet-type* > **from any** (source)
- **to** < *ip-address* [/ *mask*] | **any** > (destination)
- [*port-#*] (optional TCP or UDP application port numbers used when the packet type is TCP or UDP)

ACL: See Access Control List, below.

Access Control List (ACL): A list (or set) consisting of one or more explicitly configured Access Control Entries (ACEs) and terminating with an implicit “deny” default which drops any IP packets that do not have a match with any explicit ACE in the named ACL. An ACL can be “standard” or “extended”. See “Standard ACL” and “Extended ACL”. Both can be applied in any of the following ways:

- **Static Port ACL:** an ACL assigned to filter inbound traffic on a specific switch port

- **Dynamic Port ACL:** dynamic ACL assigned to a port by a RADIUS server to filter inbound traffic from an authenticated client on that port

An ACL can be configured on an interface as a static port ACL. (Dynamic port ACLs are configured on a RADIUS server.)

ACL Mask: Follows a destination IP address listed in an ACE. Defines which bits in a packet's corresponding IP addressing must exactly match the IP addressing in the ACE, and which bits need not match (wildcards).

DA: The acronym for *Destination IP Address*. In an IP packet, this is the destination IP address carried in the header, and identifies the destination intended by the packet's originator.

Deny: An ACE configured with this action causes the switch to drop a packet for which there is a match within an applicable ACL.

Deny Any Any: An abbreviated form of **deny in ip from any to any**, which denies any inbound IP traffic from any source to any destination.

Dynamic Port ACL: An ACL application type in which the ACL is assigned by a RADIUS server to a port to filter all inbound IP traffic from a client authenticated by the server for that port, regardless of whether the traffic is switched or routed. Filtering can be specified to include all IP traffic or specific IP applications or protocol types. Destination criteria can include a single destination IP address, a group of contiguous IP addresses, an IP subnet, or any IP destination.

Implicit Deny: If the switch finds no matches between an inbound packet and the configured criteria in an applicable ACL, then the switch denies (drops) the packet with an implicit "deny IP any/any" operation. You can preempt the implicit "deny IP any/any" in a given ACL by configuring **permit in ip from any to any** as the last explicit ACE in the ACL. Doing so permits any inbound IP packet that is not explicitly permitted or denied by other ACEs configured sequentially earlier in the ACL. Unless otherwise noted, "implicit deny IP any" refers to the "deny" action enforced by both standard and extended ACLs.

Inbound Traffic: For the purpose of defining where the switch applies ACLs to filter traffic, inbound traffic is any IP packet that *enters the switch* from a given client on a given port.

NAS (Network Attached Server): In this context, refers to a ProCurve switch configured for RADIUS operation.

Outbound Traffic: For defining the points where the switch applies an ACL to filter traffic, outbound traffic is routed traffic *leaving the switch* through a VLAN interface (or a subnet in a multinetted VLAN). “Outbound traffic” can also apply to switched traffic leaving the switch on a VLAN interface.

Permit: An ACE configured with this action allows the switch to forward an inbound packet for which there is a match within an applicable ACL.

Permit Any Any: An abbreviated form of **permit in ip from any to any**, which permits any inbound IP traffic from any source to any destination.

RADIUS-Based ACL: See “Dynamic Port ACL”.

Static Port ACL: An ACL statically configured on a specific port, group of ports, or trunk. A static port ACL filters all incoming traffic on the port, regardless of whether it is switched or routed.

VSA (Vendor-Specific-Attribute): A value used in a RADIUS-based configuration to uniquely identify a networking feature that can be applied to a port on a given vendor’s switch during an authenticated client session.

Wildcard: The part of a mask that indicates the bits in a packet’s IP addressing that do not need to match the corresponding bits specified in an ACL. See also **ACL Mask** on page 6-7.

Overview of RADIUS-Assigned, Dynamic Port ACLs

Dynamic port ACLs enhance network and switch management access security and traffic control by permitting or denying authenticated client access to specific network resources and to the switch management interface. This includes preventing clients from using TCP or UDP applications (such as Telnet, SSH, Web browser, and SNMP) if you do not want their access privileges to include these capabilities.

This feature is designed for use on the network edge to accept RADIUS-assigned, per-port ACLs (dynamic port ACLs) for Layer-3 filtering of IP traffic entering the switch from authenticated clients. A given dynamic port ACL is identified by a unique username/password pair or client MAC address, and applies only to IP traffic entering the switch from clients that authenticate with the unique credentials. The switch allows multiple dynamic port ACLs on a given port, up to the maximum number of authenticated clients allowed on the port.

A dynamic port ACL filters IP traffic entering the switch from the client whose authentication initiated the ACL assignment. Filtering criteria is based on destination and/or IP traffic type (such as TCP and UDP traffic) and traffic counter options. Implementing the feature requires:

- RADIUS authentication using the 802.1X, Web authentication, or MAC authentication services available on the switch to provide client authentication services
- configuring the ACLs on the RADIUS server (instead of the switch), and assigning each ACL to the username/password pair or MAC address of the clients you want the ACLs to support

Using RADIUS to dynamically apply per-port ACLs to edge ports enables the switch to filter IP traffic coming from outside the network, thus removing unwanted IP traffic as soon as possible and helping to improve system performance.

Note

A dynamic port ACL assignment filters all inbound IP traffic from an authenticated client on a port, regardless of whether the client's IP traffic is to be switched or routed.

Dynamic port ACLs can be used either with or without PCM and IDM support. (Refer to “Optional PCM and IDM Applications” on page 6-2.)

ACLs enhance network security by blocking selected IP traffic, and can serve as one aspect of network security. *However, because ACLs do not protect from malicious manipulation of data carried in IP packet transmissions, they should not be relied upon for a complete edge security solution.*

The ACLs described in this section do not screen non-IP traffic such as AppleTalk and IPX.

Contrasting Dynamic and Static ACLs

Table 6-1, below, highlights several key differences between the static ACLs configurable on switch ports, and the dynamic port ACLs that can be assigned to individual ports by a RADIUS server.

Table 6-1. Contrasting Dynamic and Static ACLs

Dynamic Port ACLs	Static Port ACLs
Configured in client accounts on a RADIUS server.	Configured on switch ports.
Designed for use on the edge of the network where filtering of IP traffic entering the switch from individual, authenticated clients is most important and where clients with differing access requirements are likely to use the same port.	Designed for use where the filtering needs focus on static configurations covering: <ul style="list-style-type: none">switched or routed IP traffic entering the switch from multiple sources or from unauthenticated sourcesIP traffic from multiple sources and having a destination on the switch itself
Implementation requires client authentication.	Client authentication not a factor.
Identified by the credentials (username/password pair or the MAC address) of the specific client the ACL is intended to service.	Identified by a number in the range of 1-199 or an alphanumeric name.
Supports dynamic assignment to filter only the IP traffic entering the switch from an authenticated client on the port to which the client is connected. (IP traffic can be routed or switched, and includes IP traffic having a DA on the switch itself.)	Supports static assignments to filter switched or routed IP traffic entering the switch, or routed IP traffic leaving the switch.
When the authenticated client session ends, the switch removes the RADIUS-assigned (dynamic port) ACL from the client port.	Remains statically assigned to the port.
Allows one RADIUS-assigned (dynamic port) ACL per authenticated client on a port. (Each such ACL filters traffic from a different, authenticated client.) Note: The switch provides ample resources for supporting RADIUS-assigned ACLs and other features. However, the actual number of ACLs supported depends on the switch's current feature configuration and the related resource requirements. For more information, refer to the appendix titled "Monitoring Resources" in the <i>Management and Configuration Guide</i> for your switch.	Supports static ACLs
Supports only extended ACLs. (Refer to Terminology.)	Supports standard and extended ACLs
A given dynamic port ACL filters only the IP traffic entering the switch from the authenticated client corresponding to that ACL, and does not filter IP traffic inbound from other authenticated clients. (The traffic source is not a configurable setting.)	A static port ACL applied on a port filters all traffic entering the switch through that port.

Dynamic Port ACLs	Static Port ACLs
Requires client authentication by a RADIUS server configured to dynamically assign an ACL to the client port, based on client credentials.	No client authentication requirement.
ACEs allow a counter (cnt) option that causes a counter to increment when there is a packet match.	ACEs allow a log option that generates a log message whenever there is a packet match with a “deny” ACE.

Caution Regarding the Use of Source Routing

Source routing is enabled by default on the switch and can be used to override ACLs. For this reason, if you are using ACLs to enhance network security, the recommended action is to use the **no ip source-route** command to disable source routing on the switch. (If source routing is disabled in the running-config file, the **show running** command includes “**no ip source-route**” in the running-config file listing.)

How a RADIUS Server Applies a Dynamic Port ACL to a Switch Port

A dynamic port ACL configured on a RADIUS server is identified and invoked by the unique credentials (username/password pair or a client MAC address) of the specific client the ACL is designed to service. Where the username/password pair is the selection criteria, the corresponding ACL can also be used for a group of clients that all require the same ACL policy and use the same username/password pair. Where the client MAC address is the selection criteria, only the client having that MAC address can use the corresponding ACL. When a RADIUS server authenticates a client, it also assigns the ACL configured with that client’s credentials to the port. The ACL then filters the client’s inbound IP traffic and denies (drops) any such traffic that is not explicitly permitted by the ACL. (Every ACL ends with an implicit **deny in ip from any to any** (“deny any any”) ACE that denies IP traffic not specifically permitted by the ACL.) When the client session ends, the switch removes the dynamic port ACL from the client port.

Notes

Included in any dynamic port ACL, there is an implicit **deny in ip from any to any** (“deny any any”) command that results in a default action to deny any inbound IP traffic that is not specifically permitted by the ACL. To override this default, use an explicit **permit in ip from any to any** (“permit any any”) as the last ACE in the ACL.

On a given port, dynamic port ACL filtering occurs only for the traffic entering the switch from the client whose authentication configuration on the server includes a dynamic port ACL. Traffic entering the switch from another authenticated client (on the same port) whose authentication configuration on the server does not include a dynamic port ACL will *not* be filtered by an ACL assigned to the port for any other authenticated client.

Multiple Clients Sharing the Same Dynamic Port ACL. When multiple clients supported by the same RADIUS server use the same credentials, they will all be serviced by different instances of the same ACL. (The actual IP traffic inbound from any client on the switch carries a source MAC address unique to that client. The dynamic port ACL uses this MAC address to identify the traffic to be filtered.)

Multiple ACL Application Types on an Interface. The switch allows simultaneous use of all supported ACL application types on an interface.

General ACL Features, Planning, and Configuration

These steps suggest a process for using dynamic port ACLs to establish access policies for client IP traffic.

1. Determine the policies you want to enforce for authenticated client traffic inbound on the switch.
2. Plan ACLs to execute traffic policies:
 - Apply ACLs on a per-client basis where individual clients need different traffic policies or where each client must have a different username/password pair or will authenticate using MAC authentication.
 - Apply ACLs on a client group basis where all clients in a given group can use the same traffic policy and the same username/password pair.
3. Configure the ACLs on a RADIUS server accessible to the intended clients.
4. Configure the switch to use the desired RADIUS server and to support the desired client authentication scheme. Options include 802.1X, Web authentication, or MAC authentication. (Note that the switch supports the option of simultaneously using 802.1X with either Web or MAC authentication.)

5. Test client access on the network to ensure that your RADIUS-based ACL application is properly enforcing your policies.

For further information common to all ACL applications, refer to the following sections in chapter 9, “Access Control Lists (ACLs)”:

- “Features Common to All ACLs” on page 9-10
- “General Steps for Planning and Configuring ACLs” on page 9-11
- “Planning an ACL Application” on page 9-16

The Packet-filtering Process

Packet-Filtering in an applied ACL is sequential, from the first ACE in the ACL to the implicit “deny any” following the last explicit ACE. This operation is the same regardless of whether the ACL is applied dynamically from a RADIUS server or statically in the switch configuration. For details of this process, refer to “ACL Operation” in chapter 9, “Access Control Lists (ACLs)”.

Note

If a dynamic port ACL permits an authenticated client’s inbound IP packet, but the client port is also configured with a static port ACL, then the packet will also be filtered by these other ACLs. If there is a match with a deny ACE in any of these ACLs, the switch drops the packet.

Caution

ACLs can enhance network security by blocking selected IP traffic, and can serve as one aspect of maintaining network security. *However, because ACLs do not provide user or device authentication, or protection from malicious manipulation of data carried in IP packet transmissions, they should not be relied upon for a complete security solution.*

Operating Rules for Dynamic Port ACLs

- **Relating a Client to a Dynamic Port ACL:** A dynamic port ACL for a particular client must be configured in the RADIUS server under the authentication credentials the server should expect for that client. (If the client must authenticate using 802.1X and/or Web Authentication, the username/password pair forms the credential set. If authen-

tication is through MAC Authentication, then the client MAC address forms the credential set.) For more on this topic, refer to “Configuring an ACL in a RADIUS Server” on page 6-15.

- **Multiple Clients Using the Same Username/Password Pair:** Multiple clients using the same username/password pair will use duplicate instances of the same ACL.
- **Limits for ACEs in Dynamic Port ACLs:** The switch supports up to 80 characters in a single ACE. Exceeding this limit causes the related client authentication to fail.
- **Effect of Dynamic Port ACLs on Inbound Traffic for Two Clients on the Same Port:** On a port configured for 802.1X *user-based* access where up to two clients are connected, if a given client's authentication results in a dynamic port ACL assignment, then the authentication of the other client concurrently using the port must also include a dynamic port ACL assignment. Thus, if a RADIUS server is configured to assign a dynamic port ACL when client “X” authenticates, but is not configured to do the same for client “Y”, then traffic from client “Y” will be blocked whenever client “X” is authenticated on the port (and client “Y” will be deauthenticated). For this reason, if two clients are authenticated on a port, a separate dynamic port ACL must be assigned by a RADIUS server for each authenticated client. Inbound IP traffic from a client whose authentication does not result in a dynamic port ACL assignment will be blocked and the client will be deauthenticated. Also, if 802.1X *port-based* access is configured on the port, only one client can be authenticated on the port at any given time. In this case, no other inbound client traffic is allowed.

Configuring an ACL in a RADIUS Server

This section provides general guidelines for configuring a RADIUS server to specify dynamic port ACLs. Also included is an example configuration for a FreeRADIUS server application. However, to configure support for these services on a specific RADIUS server application, please refer to the documentation provided with the application.

Elements in a Dynamic Port ACL Configuration. A dynamic port ACL configuration in a RADIUS server has the following elements:

- vendor and ACL identifiers:
 - ProCurve (HP) Vendor-Specific ID: 11
 - Vendor-Specific Attribute for ACLs: 61 (string = HP-IP-FILTER-RAW)

- Setting: HP-IP-FILTER-RAW = < “permit” or “deny” ACE >

(Note that the “string” value and the “Setting” specifier are identical.)

■ ACL configuration, including:

- one or more explicit “permit” and/or “deny” ACEs created by the system operator
- implicit deny any any ACE automatically active after the last operator-created ACE

Example of Configuring a Dynamic Port ACL Using the FreeRADIUS Application. This example illustrates one method for configuring dynamic port ACL support for two different client identification methods (username/password and MAC address). For information on how to configure this functionality on other RADIUS server types, refer to the documentation provided with the server.

1. Enter the ProCurve vendor-specific ID and the ACL VSA in the FreeRADIUS dictionary file:




Figure 6-3. Example of Configuring the VSA for Dynamic Port ACLs in a FreeRADIUS Server

2. Enter the switch IP address, NAS (Network Attached Server) type, and the key in the FreeRADIUS **clients.conf** file. For example, if the switch IP address is 10.10.10.125 and the key is “1234”, you would enter the following in the server’s **clients.conf** file:




Figure 6-4. Example of Configuring the Switch’s Identity Information in a FreeRADIUS Server

3. For a given client username/password pair or MAC address, create an ACL by entering one or more ACEs in the FreeRADIUS “users” file. Enter the ACEs in an order that promotes optimum traffic management and conser-

vation of system resources, and remember that every ACL you create automatically includes an implicit **deny in ip from any to any** ACE. For example, suppose that you wanted to create identical ACL support for the following:

- a client having a username of “mobile011” and a password of “run101112”
- a client having a MAC address of 08 E9 9C 4F 00 19

The ACL in this example must achieve the following:

- permit http (TCP port 80) traffic from the client to the device at 10.10.10.101
- deny http (TCP port 80) traffic from the client to all other devices
- permit all other traffic from the client to all other devices

To configure the above ACL, you would enter the username/password and ACE information shown in figure6-5 into the FreeRADIUS **users** file.

Note

For syntax details on dynamic port ACLs, refer to the next section, “Format Details for ACEs Configured in a Dynamic Port ACL”.

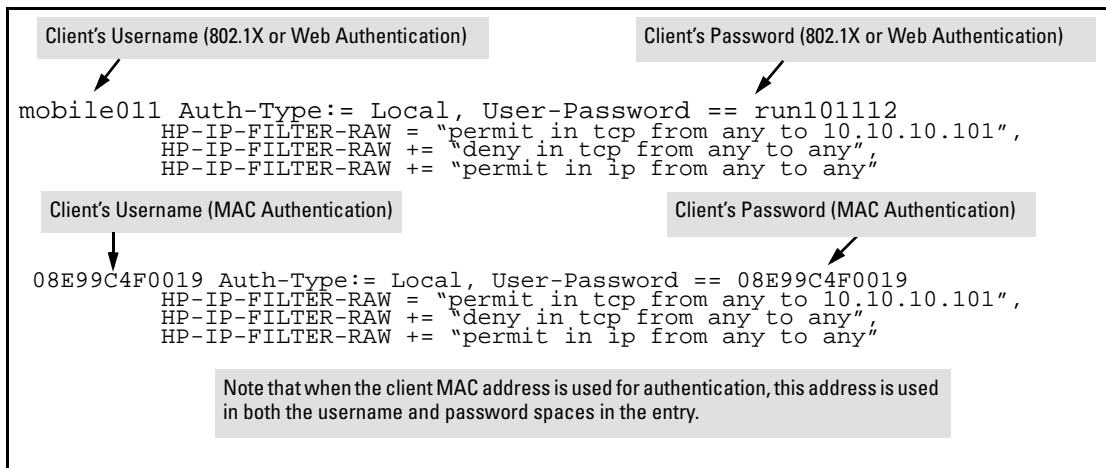


Figure 6-5. Example of Configuring the FreeRADIUS Server To Support ACLs for the Indicated Clients

Format Details for ACEs Configured in a Dynamic Port ACL.

Any instance of a dynamic port ACL is structured to filter authenticated client traffic as follows:

- Applies only to inbound client traffic on the switch port the authenticated client is using.
- Allows only the “any” source address (for any authenticated IP device connected to the port).
- Applies to all IP traffic from the authenticated client or to a specific type of IP traffic type from the client. Options include TCP, UDP, or any other type of IP traffic that is identified by an IP protocol number. (More information on protocol numbers is provided in the following ACL syntax description.) Has one of the following destination types:
 - A specific IP address
 - A contiguous series of IP address or an entire subnet
 - Any IP address
- Where the traffic type is either TCP or UDP, the ACE can optionally include one or more TCP or UDP port numbers.

Configuring ACE Syntax in RADIUS Servers

The following syntax and operating information applies to ACLs configured in a RADIUS server.

ACE Syntax: < permit | deny > in < ip | *ip-protocol-value* > from any to < *ip-addr* > [< *mask* >] | > [*tcp/udp-ports*] [cnt]

< permit | deny >: *Specifies whether to forward or drop the identified IP traffic type from the authenticated client. (For information on explicitly permitting or denying all inbound IP traffic from an authenticated client, or for implicitly denying all such IP traffic not already permitted or denied, refer to “Configuration Notes” on page 6-19.)*

in: *Required keyword specifying that the ACL applies only to the traffic inbound from the authenticated client.*

< ip | *ip-protocol-value* >: *Options for specifying the type of traffic to filter.*

ip: *This option applies the ACL to all IP traffic from the authenticated client.*

ip-protocol-value: *This option applies the ACL to the type of IP traffic specified by either a protocol number or by tcp or udp. The range of protocol numbers is 0-255, and you can substitute 6 for TCP or 17 for UDP. (Protocol numbers are defined in RFC 2780. For a complete listing, refer to “Protocol Numbers” under “Protocol Number Assignment Services” on the Web site of the Internet Assigned Numbers Authority at www.iana.com.) Some examples of protocol numbers include:*

1 = ICMP	17 = UDP
2 = IGMP	41 = IPv6
6 = TCP	

from any: *Required keywords specifying the (authenticated) client source. (Note that a dynamic port ACL assigned to a port filters only the inbound traffic having a source MAC address that matches the MAC address of the client whose authentication invoked the ACL assignment.)*

to: *Required destination keyword.*

< ip-addr >: *Specifies a single destination IP address.*

< ip-addr /< mask >: *Specifies a series of contiguous destination IP addresses or all destination IP addresses in a subnet. The < mask > is CIDR notation for the number of leftmost bits in a packet's destination IP address that must match the corresponding bits in the destination IP address listed in the ACE. For example, a destination of 10.100.17.1/24 in the ACE means that a match occurs when an inbound packet (of the designated IP type) from the authenticated client has a destination IP address where the first three octets are 10.100.17. (The fourth octet is a wildcard, and can be any value up to 255.)*

any: *Specifies any IP destination address. Use this option when you want the ACL action to apply to all traffic of the designated type, regardless of destination.*

[tcp/udp-ports]: *Optional TCP or UDP port specifier. Used when the ACL is intended to filter client TCP or UDP traffic with one or more specific TCP or UDP destination port numbers. You can specify port numbers as individual values and/or ranges. For example, the following ACE denies any UDP traffic from an authenticated client that has a DA of any IP address and a UDP destination port of 135, 137-139, or 445:*

deny in udp from any to any 135, 137-139, 445.

[cnt]: *Optional counter specifier for a dynamic port ACL. When used in an ACL, the counter increments each time there is a “match” with a permit or deny ACE. This option requires that you configure the switch for RADIUS accounting.*

Configuration Notes

Explicitly Permitting Any IP Traffic. Entering a **permit in ip from any to any** (permit any any) ACE in an ACL permits all IP traffic not previously permitted or denied by that ACL. Any ACEs listed after that point do not have any effect.

Explicitly Denying Any IP Traffic. Entering a **deny in ip from any to any** ACE in an ACL denies all IP traffic not previously permitted or denied by that ACL. Any ACEs listed after that point have no effect.

Implicitly Denying Any IP Traffic. For any packet being filtered by a static port ACL, there will always be a match. That is, any packet that does not have a match with an explicit permit or deny ACE in the list will match with the implicit **deny in ip from any to any** that is automatically implied at the end of the list. Thus, the ACL denies any IP packet it filters that does not match any explicitly configured ACE. If you want an ACL to permit any packets that

are not explicitly denied, you must configure **permit in ip from any to any** as the last explicit ACE in the ACL. This pre-empts the implicit **deny in ip from any to any** ACE and permits packets not explicitly permitted or denied by earlier ACEs in the list.

Configuring the Switch To Support Dynamic Port ACLs

An ACL configured in a RADIUS server is identified by the authentication credentials of the client or group of clients the ACL is designed to support. When a client authenticates with credentials associated with a particular ACL, the switch applies that ACL to the switch port the client is using. To enable the switch to forward a client's credentials to the RADIUS server, you must first configure RADIUS operation and an authentication method on the switch.

1. Configure RADIUS operation on the switch:

Syntax: radius-server host < ip-address > key < key-string >

This command configures the IP address and encryption key of a RADIUS server. The server should be accessible to the switch and configured to support authentication requests from clients using the switch to access the network. For more on RADIUS configuration, refer to chapter 5, "RADIUS Authentication and Accounting".

2. Configure RADIUS network accounting on the switch (optional). RADIUS network accounting is necessary to retrieve counter information if the **cnt** (counter) option is included in any of the ACEs configured on the RADIUS server.

Syntax: aaa accounting network < start-stop | stop-only > radius

Note

Refer to the documentation provided with your RADIUS server for information on how the server receives and manages network accounting information, and how to perform any configuration steps necessary to enable the server to support network accounting data from the switch.

3. Configure an authentication method. Options include 802.1X, Web authentication, and MAC authentication. (You can configure 802.1X and either Web or MAC authentication to operate simultaneously on the same ports.)

802.1X Option:

Syntax: aaa port-access authenticator < port-list >
aaa authentication port-access chap-radius
aaa port-access authenticator active

These commands configure 802.1X port-based access control on the switch, and activates this feature on the specified ports. For more on 802.1X configuration and operation, refer to chapter 11, “Configuring Port-Based and User-Based Access Control (802.1X)” in this guide.

MAC Authentication Option:

Syntax: aaa port-access mac-based < port-list >

This command configures MAC authentication on the switch and activates this feature on the specified ports. For more on MAC authentication, refer to chapter 3, “Web and MAC Authentication”.

Web Authentication Option:

Syntax: aaa port-access web-based < port-list >

This command configures Web authentication on the switch and activates this feature on the specified ports. For more on Web authentication, refer to chapter 3, “Web and MAC Authentication”.

Displaying the Current Dynamic Port ACL Activity on the Switch

These commands output data indicating the current ACL activity imposed per-port by RADIUS server responses to client authentication.

Syntax: show access-list radius < port-list >

*For the specified ports, this command lists the explicit ACEs, switch port, and client MAC address for each ACL dynamically assigned by a RADIUS server as a response to client authentication. If **cnt** (counter) is included in an ACE, then the output includes the current number of inbound packet matches the switch has detected in the current session for that ACE.*

Note: *If a client authenticates but the server does not return a dynamic port ACL to the client port, then the server does not have a valid ACL configured and assigned to that client’s authentication credentials.*

For example, the following output shows that a RADIUS server has assigned an ACL to port B1 to filter inbound traffic from an authenticated client identified by a MAC address of 00-11-85-C6-54-7D.

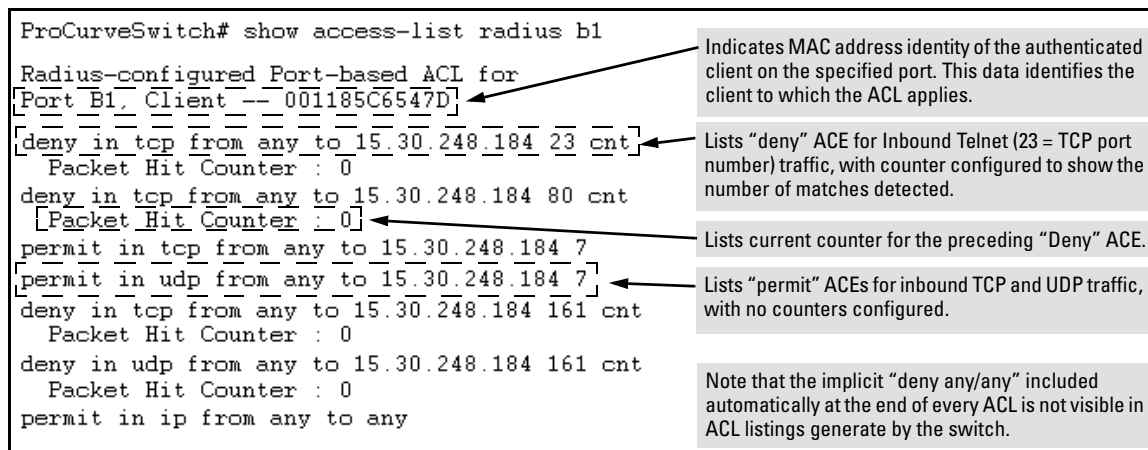


Figure 6-6. Example Showing a Dynamic Port ACL Application to a Currently Active Client Session

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

For ports, in < **port-list** > that are configured for authentication, this command indicates whether there are any RADIUS-assigned features active on the port(s). (Any ports in < **port-list** > that are not configured for authentication do not appear in this listing.)

Port: Port number of port configured for authentication.

Status: Port connection status:

Open = active connection with an external device

Closed = no active connection with an external device

Current VLAN ID: VLAN ID (VID) of the VLAN currently supporting the active connection.

Current Port CoS: Indicates the status of the current 802.1p priority setting for inbound traffic.

No-override: Indicates that no RADIUS-assigned 802.1p priority is currently active on the indicated port. (For more on traffic prioritization for the switches covered in this guide, refer to the chapter titled "Quality of Service (QoS): Managing Bandwidth More Effectively", in this guide.)

0 - 7: Indicates that the displayed 802.1p priority has been assigned by a RADIUS server to inbound traffic on the indicated port for a currently active, authenticated client session. This assignment remains active until the session ends.

RADIUS ACL Applied?: Indicates whether a dynamic port ACL is currently active on the port.

Yes: *An ACL has been assigned by a RADIUS server to inbound traffic on the indicated port for a currently active, authenticated client session. This assignment remains active until the session ends.*

No: *There is no dynamic port ACL currently active on the indicated port.*

```
ProCurve (config)# show port-access authenticator 1-10
Port Access Authenticator Status
Port-access authenticator activated [No] : No
Allow RADIUS-assigned dynamic (GVRP) VLANs [No] : No

Port  Status      Current  Current  RADIUS ACL
-----  -----  VLAN ID  Port COS  Applied?
1       Open       1        7         Yes
2       Closed    1        No-override No
3       Open       1        No-override Yes
```

Figure 6-7. Example of Output Showing Current RADIUS-Applied Features

Event Log Messages

Message	Meaning
ACE parsing error, permit/deny keyword <ace-#> client <mac-address> port <port-#>.	Notifies of a problem with the permit/deny keyword in the indicated ACE included in the access list for the indicated client on the indicated switch port.
Could not add ACL entry.	Notifies that the ACE entry could not be added to the internal ACL storage.
Could not create ACL entry.	Notifies that the ACL could not be added to the internal ACL storage.
Could not add ACL, client mac<mac-address>port <port-#>, at max per-port ACL quantity.	Notifies that the ACL could not be added because the per-port ACL quantity would be exceeded.
ACE parsing error, IN keyword, <ace-#> client <mac-address> port <port-#>.	Notifies of a problem with the IN keyword in the indicated ACE of the access list for the indicated client on the indicated switch port.
ACE parsing error, protocol field, <ace-#> client <mac-address> port <port-#>.	Notifies of a problem with the protocol field in the indicated ACE of the access list for the indicated client on the indicated switch port.
ACE parsing error, FROM keyword, <ace-#> client <mac-address> port <port-#>.	Notifies of a problem with the FROM keyword in the indicated ACE of the access list for the indicated client on the indicated switch port.
ACE parsing error, ANY keyword, <ace-#> client <mac-address> port <port-#>.	Notifies of a problem with the ANY keyword in the indicated ACE of the access list for the indicated client on the indicated switch port.
ACE parsing error, TO keyword, <ace-#> client <mac-address> port <port-#>.	Notifies of a problem with the TO keyword in the indicated ACE of the access list for the indicated client on the indicated switch port.
ACE parsing error, destination IP, <ace-#> client <mac-address> port <port-#>.	Notifies of a problem with the destination IP field in the indicated ACE of the access list for the indicated client on the indicated switch port.
ACE parsing error, tcp/udp ports, <ace-#> client <mac-address> port <port-#>.	Notifies of a problem with the TCP/UDP port field in the indicated ACE of the access list for the indicated client on the indicated switch port.
Rule limit per ACL exceeded. <ace-#> client <mac-address> port <port-#>.	Notifies that an ACL has too many rules.
Duplicate mac. An ACL exists for client. Deauthenticating second. client <mac-address> port <port-#>.	Notifies that an ACL for this mac on this port already exists.

Message	Meaning
Invalid Access-list entry length, client <mac-address> port <port-#>.	Notifies that the string configured for an ACE entry on the Radius server exceeds 80 characters.
Memory allocation failure for IDM ACL.	Notifies of a memory allocation failure for a dynamic port ACL assigned by a RADIUS server performing client authentication. (This message is used in IDM and non-IDM environments.)

Causes of Client Deauthentication Immediately After Authenticating

- ACE formatted incorrectly in the RADIUS server
 - “from”, “any”, or “to” keyword missing
 - An IP protocol number in the ACE exceeds 255.
 - An optional UDP or TCP port number is invalid, or a UDP/TCP port number is specified when the protocol is neither UDP or TCP.
- A dynamic port ACL limit has been exceeded.
 - An ACE in the ACL for a given authenticated client exceeds 80 characters.
 - The TCP/UDP port-range quantity of 14 per slot or port group has been exceeded.

Monitoring Shared Resources

Currently active, RADIUS-based authentication sessions (including ProCurve IDM client sessions) using dynamic port ACLs share internal routing switch resources with several other features. The routing switch provides ample resources for all features. However, if the internal resources do become fully subscribed, new RADIUS-based sessions using dynamic port ACLs cannot be authenticated until the necessary resources are released from other applications.

Configuring Secure Shell (SSH)

Contents

Overview	7-2
Terminology	7-4
Prerequisite for Using SSH	7-5
Public Key Formats	7-5
Steps for Configuring and Using SSH for Switch and Client Authentication	7-6
General Operating Rules and Notes	7-8
Configuring the Switch for SSH Operation	7-9
1. Assign Local Login (Operator) and Enable (Manager) Password .	7-9
2. Generate the Switch's Public and Private Key Pair	7-10
3. Provide the Switch's Public Key to Clients	7-12
4. Enable SSH on the Switch and Anticipate SSH Client Contact Behavior	7-15
5. Configure the Switch for SSH Authentication	7-18
6. Use an SSH Client To Access the Switch	7-21
Further Information on SSH Client Public-Key Authentication	7-22
Messages Related to SSH Operation	7-28

Overview

Feature	Default	Menu	CLI	Web
Generating a public/private key pair on the switch	No	n/a	page 7-10	n/a
Using the switch's public key	n/a	n/a	page 7-12	n/a
Enabling SSH	Disabled	n/a	page 7-15	n/a
Enabling client public-key authentication	Disabled	n/a	pages 7-19, 7-22	n/a
Enabling user authentication	Disabled	n/a	page 7-18	n/a

The ProCurve switches covered in this guide use Secure Shell version 1 or 2 (SSHv1 or 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.)

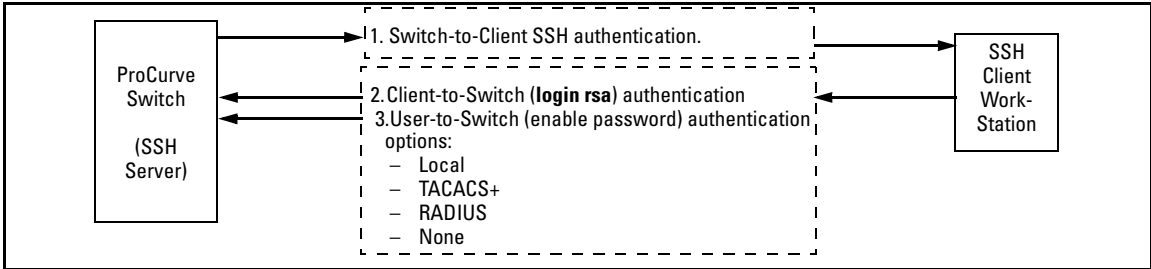


Figure 7-1. Client Public Key Authentication Model

Note

SSH in the ProCurve is based on the OpenSSH software toolkit. For more information on OpenSSH, visit www.openssh.com.

Switch SSH and User Password Authentication . This option is a subset of the client public-key authentication show in figure 7-1. 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 7-1, 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.

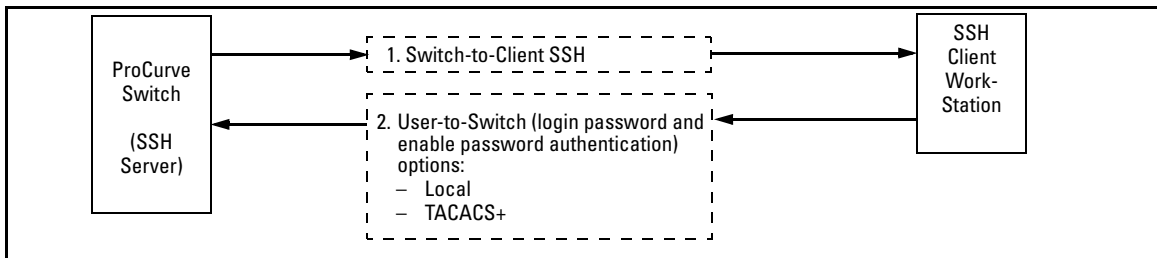


Figure 7-2. Switch/User Authentication

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 (v1/v2 shared host key & v1 server 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 figures 7-3 and 7-4 for examples of PEM-encoded ASCII and non-encoded ASCII keys.
- **Private Key:** An internally generated key used in the authentication process. A private key generated by the switch is not accessible for viewing or copying. A private key generated by an SSH client application is typically stored in a file on the client device and, together with its public key counterpart, can be copied and stored on multiple devices.
- **Public Key:** An internally generated counterpart to a private key. A device's public key is used to authenticate the device to other devices.
- **Enable Level:** Manager privileges on the switch.
- **Login Level:** Operator privileges on the switch.
- **Local password or username:** A Manager-level or Operator-level password configured in the switch.
- **SSH Enabled:** (1) A public/private key pair has been generated on the switch (**crypto key generate 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 7-10 and “4. Enable SSH on the Switch and Anticipate SSH Client Contact Behavior” on page 7-15.)

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 7-2), 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"A1B3Nz1y2+orEhYL . . . Q8D8qDH1ozu1c="*** End of Pub Key ***"
```

Comment describing public key identity.

Beginning of actual SSHv2 public key in PEM-Encoded ASCII format.

Figure 7-3. Example of Public Key in PEM-Encoded ASCII Format Common for SSHv2 Clients

```
512 37 78193303392019545793321845914508115859448079486918367079008218589443776362026267. . .
```

Bit Size

Exponent <e>

Modulus <n>

Figure 7-4. Example of Public Key in Non-Encoded ASCII Format (Common for SSHv1 Client Applications)

Steps for Configuring and Using SSH for Switch and Client Authentication

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

Table 7-1. SSH Options

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

¹ For **ssh login 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 7-23.

B. Switch Preparation

1. Assign a login (Operator) and enable (Manager) password on the switch (page 7-9).
2. Generate a public/private key pair on the switch (page 7-10).

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 7-12).
4. Enable SSH on the switch (page 7-15).
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
 - Option B:

Primary: Client public-key authentication (**login public-key** — page 7-22)
Secondary: Local password or none

Note that if you want the switch to perform client public-key authentication, you must configure the switch with Option B.
- SSH Enable (Manager) options:

Primary: Local, TACACS+, or RADIUS
Secondary: Local password or none
6. Use your SSH client to access the switch using the switch's IP address or DNS name (if allowed by your SSH client application). Refer to the documentation provided with the client application.

General Operating Rules and Notes

- Public keys generated on an SSH client must be exportable to the switch. The switch can only store ten keys 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 is not secure.
- The switch does not support outbound SSH sessions. Thus, if you Telnet from an SSH-secure switch to another SSH-secure switch, *the session is not secure*.

Configuring the Switch for SSH Operation

SSH-Related Commands in This Section	Page
show ip ssh	7-17
show crypto client-public-key [keylist-str] [< babble fingerprint >]	7-25
show crypto host-public-key [< babble fingerprint >]	7-14
show authentication	7-21
crypto key < generate zeroize > ssh [rsa]	7-11
ip ssh	7-16
key-size < 512 768 1024 >	7-16
port < 1 - 65535 default >	7-16
timeout < 5 - 120 >	7-16
version < 1 2 1-or-2 >	7-16
aaa authentication ssh	
login < local tacacs radius public-key >	7-18, 7-20
< local none >	7-18
enable < tacacs radius local >	7-18
< local none >	7-18
copy tftp pub-key-file <tftp server IP> <public key file>	7-25
clear crypto client-public-key [keylist-str]	7-26

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 7-5. 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 7-14.)

`[fingerprint]`

Displays fingerprints of the switch's public key in hexadecimal format. (See "Displaying the Public Key" on page 7-14.)

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 AAAAB3NzaClyc2EAAAABIAAAHEAnAAApdhq13Jynrs7j4lDUm8ivVm8ld2mZU5e+YZWp/T6
QzP2RsDDMZLbAHHIBrxPLjW/bRogpYD0lWuVOhTojEVjqeVuXbwmdDnyOgBc06olePwdrbQ+FZevERiA
JYG3C8NCzCRD/djXeI7FmRps8w==
-----
```

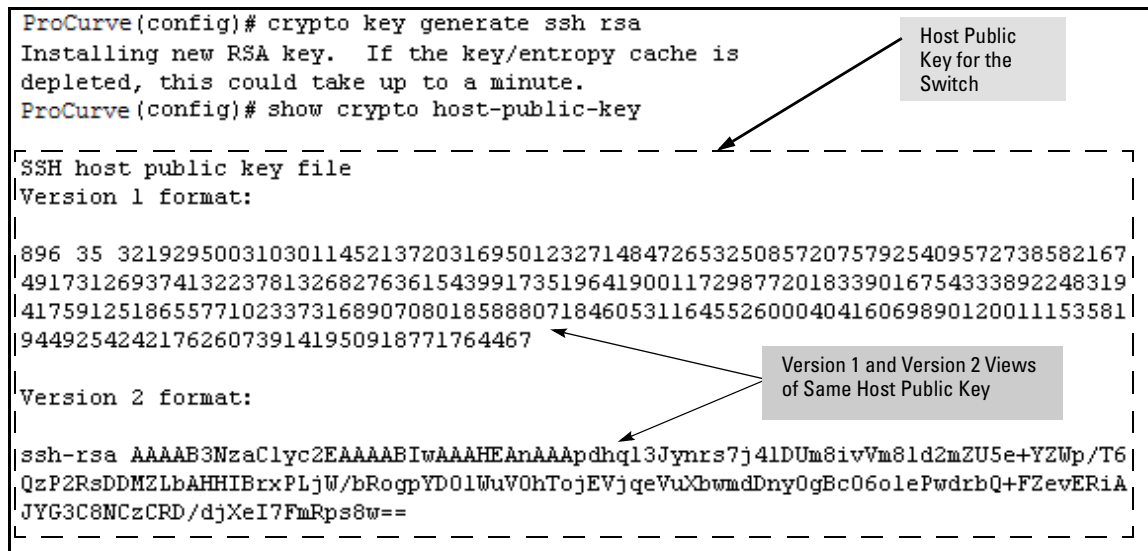


Figure 7-6. Example of Generating a Public/Private Host Key Pair for the Switch

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 7-7. 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 7-6).
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 7-8. 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 7-9. 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 7-8 on page 7-13.)
- **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 7-8 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
```

Phonetic "Hash" of Switch's Public Key

Hexadecimal "Fingerprints" of the Same Switch

Figure 7-10. Examples of Visual Phonetic and Hexadecimal Conversions of the Switch's Public Key

The two commands shown in figure 7-10 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 7-10.

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 7-18.

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 7-10.)
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 7-11).

Syntax: [no] ip ssh

Enables or disables SSH on the switch.

[key-size < 512 | 768 | 1024 >] Version 1 only

The size of the internal, automatically generated key the switch uses for negotiations with an SSH client. A larger key provides greater security; a smaller key results in faster authentication (default: 512 bits).

[port < 1-65535 | default >]

*The TCP port number for SSH connections (default: 22). **Important:** See “Note on Port Number” on page 7-17.*

[timeout < 5 - 120 >]

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

[version < 1 | 2 | 1-or-2 >]

*The version of SSH to accept connections from.
(default: 1-or-2)*

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 7-18.

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 7-11. Example of Enabling IP SSH and Listing the SSH Configuration and Status

7-17

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 8, “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 7-19) 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 7-22

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 7-9.

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**.*

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

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

Option B: Configuring the Switch for Client Public-Key SSH

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

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

(For more on these topics, refer to “Further Information on SSH Client Public-Key Authentication” on page 7-22.)

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**).*

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**.*

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:

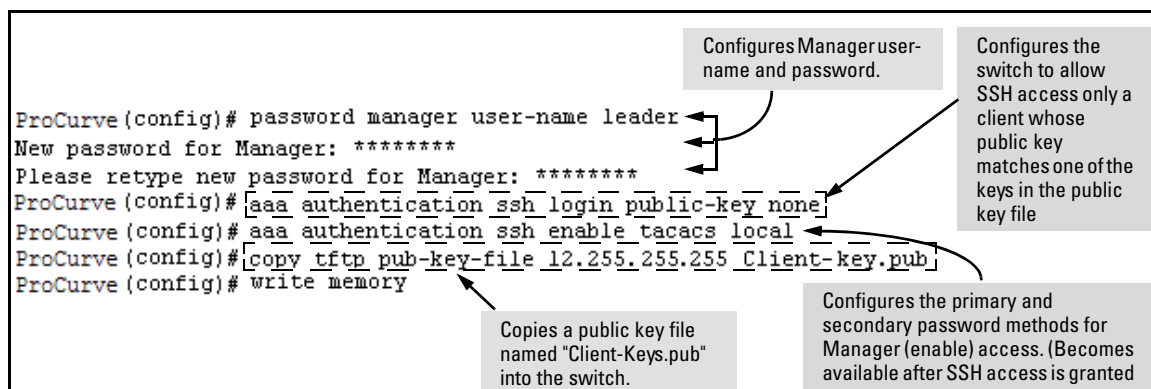


Figure 7-12. Configuring for SSH Access Requiring a Client Public-Key Match and Manager Passwords

Figure 7-13 shows how to check the results of the above commands.

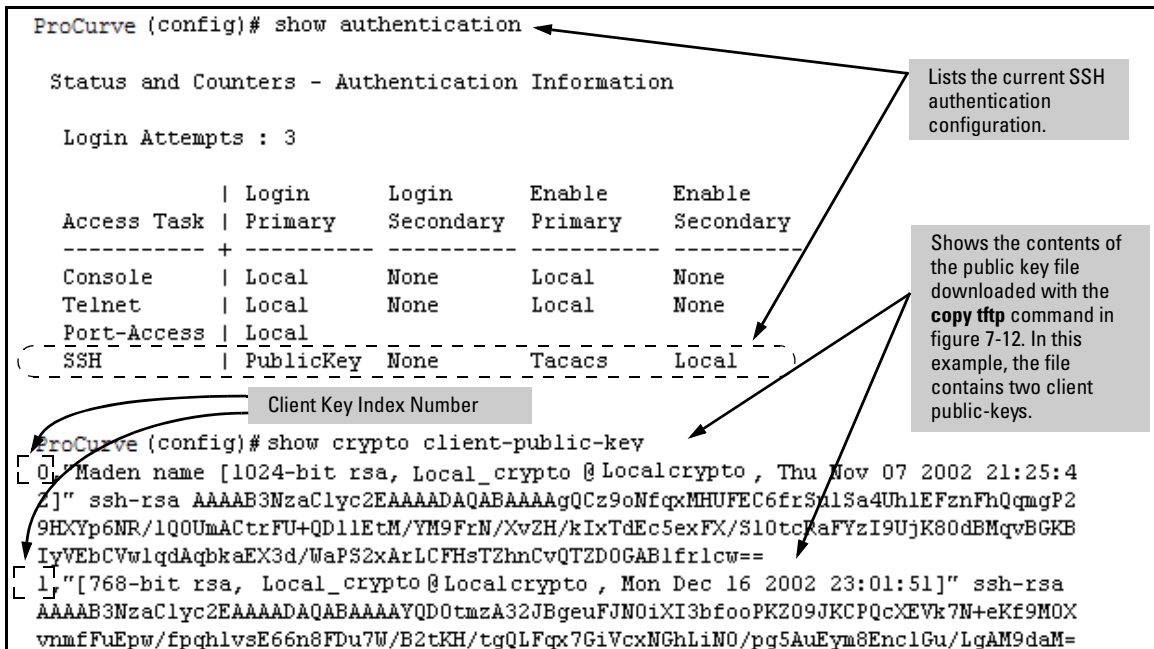


Figure 7-13. SSH Configuration and Client-Public-Key Listing From Figure 7-12

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 7-18 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 7-18.)

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.

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 7-14. Example of a Client Public Key

Notes

Comments in public key files, such as **smith@support.cairns.com** in figure 7-14, may appear in a SSH client application's generated public key. While such comments may help to distinguish one key from another, they do not pose any restriction on the use of a key by multiple clients and/or users.

Public key illustrations such as the key shown in figure 7-14 usually include line breaks as a method for showing the whole key. However, in practice, line breaks in a public key will cause errors resulting in authentication failure.

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

Property	Supported Value	Comments
Key Format	ASCII	See figure 7-8 on page 7-13. 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 7-14 on page 7-23.)

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

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

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

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

*The **babble** option converts the key data to phonetic hashes that are easier for visual comparisons.*

*The **fingerprint** option converts the key data to phonetic 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 AAAAB3NzaClyc2EAAAADAQABAAQgQCz9oNfqxMHUFEC6frSulSa4Uh1EFznFhQqmgP2
9HXVp6NR/1QOUmACtrFU+QD11EtM/YM9FrN/XvZH/kIxTdEc5exFX/S10tcRaFYzI9UjK80dBMqvBGKB
IyVEbCVwlqdAqbkEX3d/WaPS2xArLCFHsTZhnCvQTZD0GABlfrlcw==
1."[768-bit rsa, Jamie_wilson@Jamiewilson, Mon Dec 16 2002 23:01:51]" ssh-rsa
AAAAB3NzaClyc2EAAAADAQABAAQgYQD0tmzA32JBgeuFJN0iXI3bfooPKZ09JKCPQcXEVk7N+eKf9M0X
vnmfFuBpw/fpqhlvsE66n8FDu7W/B2tKH/tqQLFqx7GiVcxNGhLiN0/pq5AuEym8EnclGu/LgAM9daM=
```



Figure 7-15. 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 one of the following:

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

Syntax: aaa authentication ssh login 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.

aaa authentication ssh login public-key local

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

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

Caution

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

Messages Related to SSH Operation

Message	Meaning
00000K Peer unreachable.	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 7-17.
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: overlenght 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 23 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.

Message	Meaning
Generating new RSA host key. If the cache is depleted, this could take up to two minutes.	After you execute the <code>crypto key generate ssh [rsa]</code> command, the switch displays this message while it is generating the key.
Host RSA key file corrupt or not found. Use 'crypto key generate 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.

Configuring Secure Socket Layer (SSL)

Contents

Overview	8-2
Terminology	8-3
Prerequisite for Using SSL	8-5
Steps for Configuring and Using SSL for Switch and Client Authentication	8-5
General Operating Rules and Notes	8-6
Configuring the Switch for SSL Operation	8-7
1. Assign Local Login (Operator) and Enable (Manager) Password .	8-7
2. Generate the Switch's Server Host Certificate	8-9
3. Enable SSL on the Switch and Anticipate SSL Browser Contact Behavior	8-17
Common Errors in SSL Setup	8-21

Overview

Feature	Default	Menu	CLI	Web
Generating a Self Signed Certificate on the switch	No	n/a	page 8-9	page 8-13
Generating a Certificate Request on the switch	No	n/a	n/a	page 8-15
Enabling SSL	Disabled	n/a	page 8-17	page 8-19

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 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 8-1, 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 8-1. 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 8-7).
2. Generate a host certificate on the switch (page 8-9).
 - 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 8-17).
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 8-19
show config	page 8-19
show crypto host-cert	page 8-12
crypto key	
generate cert [rsa] <512 768 1024>	page 8-10
zeroize cert	page 8-10
crypto host-cert	
generate self-signed [arg-list]	page 8-10
zeroize	page 8-10

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.

The screenshot shows the HP ProCurve Switch web interface. At the top, there is a header bar with "HP ProCurve Switch" and "Status: Information". Below this is a navigation bar with tabs: Identity, Status, Configuration, Security, Diagnostics, and Support. The Security tab is selected. Below the Security tab is a sub-navigation bar with tabs: Device Passwords, Authorized Addresses, Port Security, Intrusion Log, and SSL. The Device Passwords sub-tab is selected. The main content area is divided into two sections: Read-Only Access and Read-Write Access. Each section contains three input fields: User Name, Password, and Confirm Password. At the bottom right, there are two buttons: Apply Changes and Clear Changes. An arrow points to the Security tab, and another arrow points to the Device Passwords sub-tab. A label "Password Button" with an arrow points to the Password input field in the Read-Only Access section.

Figure 8-2. 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 there certificate folder on the first use. This method does allow for a security breach on the first access to the switch. (Refer to the documentation for your browser application.)

There are two types of certificated that can be used for the switch's host certificate. The first type is a self-signed certificate, which is generated and digitally signed by the switch. Since self-signed certificates are not signed by a third-party certificate authority, there is no audit trail to a root CA certificate and no fool-proof means of verifying authenticity of certificate. The second type is a certificate authority-signed certificate, which is digitally signed by a certificate authority, has an audit trail to a root CA certificate, and can be verified unequivocally

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 8-9.)

`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 8-1, “Certificate Field Descriptions” describes these arguments.

Table 8-1. 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
```

Generate New Key

Generate New Certificate

Enter certificate Arguments

Figure 8-3. Example of Generating a Self-Signed Server Host certificate on the CLI for the Switch.

Note

“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 8-4. 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 8-11.)
- 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)

Configuring the Switch for SSL Operation

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

ProCurve Switch Status: Information

Identity Status Configuration **Security** Diagnostics Support

Device Passwords Authorized Addresses Port Security Intrusion Log **SSL**

SSL Settings

SSL Enable: ☒ Off Port: 443

☒ Create Certificate/ Certificate Request ☐ Use Installed Certificate

Certificate Type: Self Signed

RSA Key Size: 512

Certificate Information Fields

Validity Start Date: Month Day Year

Validity End Date: Month Day Year

Common Name: 10.255.255.255

Organization Name: Company Name

Organization Unit: Department Name

City: City

State: State

Country: US - United States

Certificate Arguments

Validity End Date:

Common Name :

Organization Name :

Organization Unit :

City :

State:

Country :

Fingerprint MD5:

SHA :

Apply Changes Clear Changes

Figure 8-5. Self-Signed Certificate generation via SSL Web Browser Interface Screen

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

1. Proceed to the **Security** tab
2. Then the **[SSL]** button

ProCurve Switch Status: Information

Identity Status Configuration **Security** Diagnostics Support

Device Passwords Authorized Addresses Port Security Intrusion Log **SSL**

SSL Settings

Current SSL Host Certificate

SSL Enable: Port:

☐ Create Certificate/ Certificate Request

Certificate Type:

RSA Key Size:

Certificate Information Fields

Validity Start Date:

Validity End Date:

Common Name:

Organization Name:

Organization Unit:

City:

State:

Country:

☒ Use Installed Certificate

Installed Certificate

Certificate Type : Self-Signed

RSA Key Size : 512 bits

Validity Start Date: 1/1/2002

Validity End Date: 1/1/2003

Common Name : 10.255.255.255

Organization Name : Hewlett Packard

Organization Unit : ProCurve Network

City : Roseville

State: Ca

Country : US

Fingerprint
MD5: BE01 E39E D49C 2575
200B 30E6 E080 38C3
CE94 BFD8 86F8 1887

SHA : BE24 F173 55D4 BE0A
4E05 2C40

Figure 8-6. 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 8-11.)
- 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.

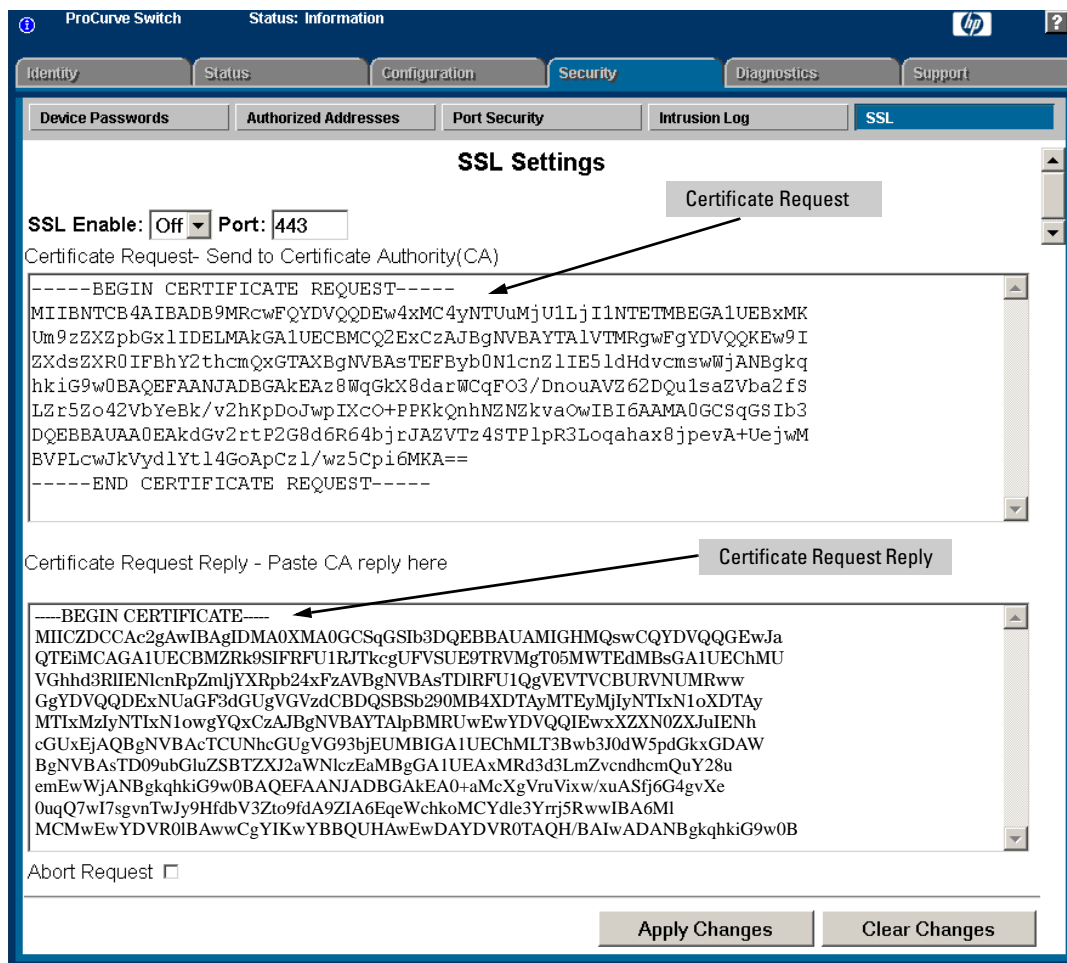


Figure 8-7. Example of a Certificate Request and Reply

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 8-9.

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 8-20.*

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 8-9.)
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 8-10).

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.

Configuring Secure Socket Layer (SSL)

Configuring the Switch for SSL Operation

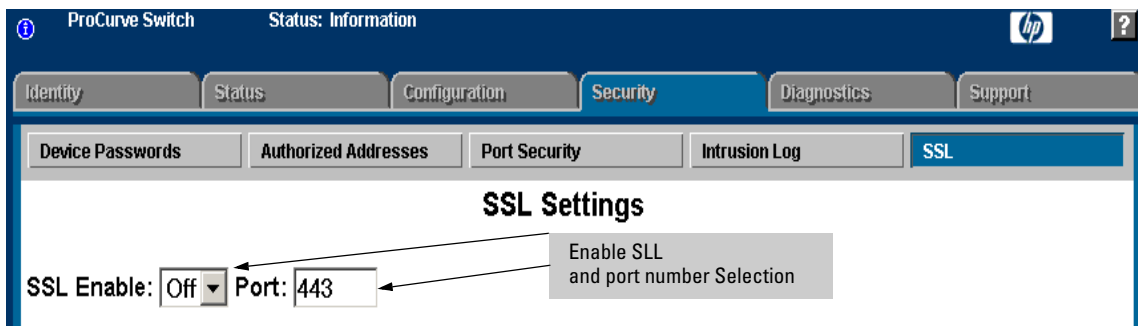


Figure 8-8. Using the web browser interface to enable SSL and select TCP port number

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 8-10.)
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 8-13.)</p> <p>You may be using a reserved TCP port. (Refer to "Note on Port Number" on page 8-20.)</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 8-17.)</p> <p>Your browser may not support SSLv3 or TLSv1 or it may be disabled. (Refer to the documentation provided for your browser.)</p>

Access Control Lists (ACLs)

Contents

Introduction	9-3
ACL Applications	9-3
Optional Network Management Applications	9-3
Optional PCM and IDM Applications	9-4
General Application Options	9-4
Terminology	9-6
Overview	9-9
Types of IP ACLs	9-9
ACL Inbound Application Points	9-9
Features Common to All ACLs	9-10
General Steps for Planning and Configuring ACLs	9-11
ACL Operation	9-12
Introduction	9-12
The Packet-Filtering Process	9-13
Planning an ACL Application	9-16
Switch Resource Usage	9-16
Managing ACL Resource Consumption	9-18
Traffic Management and Improved Network Performance	9-22
Security	9-22
Guidelines for Planning the Structure of an ACL	9-23
ACL Configuration and Operating Rules	9-24
How an ACE Uses a Mask To Screen Packets for Matches	9-25
Configuring and Assigning an ACL	9-32
Overview	9-32
ACL Configuration Structure	9-33
ACL Configuration Factors	9-36
Using the CLI To Create an ACL	9-38

Configuring and Assigning a Numbered, Standard ACL	9-39
Configuring and Assigning a Numbered, Extended ACL	9-44
Configuring a Named ACL	9-50
Enabling or Disabling ACL Filtering on an Interface	9-52
Deleting an ACL from the Switch	9-53
Displaying ACL Data	9-54
Display an ACL Summary	9-54
Display the Content of All ACLs on the Switch	9-55
Display the ACL Assignments for an Interface	9-56
Displaying the Content of a Specific ACL	9-57
Displaying the Current ACL Resources	9-59
Display All ACLs and Their Assignments in the Switch Startup-Config File and Running-Config File	9-60
Editing ACLs and Creating an ACL Offline	9-60
Using the CLI To Edit ACLs	9-60
Working Offline To Create or Edit an ACL	9-63
Enable ACL “Deny” Logging	9-67
Requirements for Using ACL Logging	9-67
ACL Logging Operation	9-67
Enabling ACL Logging on the Switch	9-68
Operating Notes for ACL Logging	9-70
General ACL Operating Notes	9-71

Introduction

Feature	Default	Menu	CLI	Web
Numbered ACLs				
Standard ACLs	None	—	9-39	—
Extended ACLs	None	—	9-44	—
Named ACLs		—	9-50	—
Enable or Disable an ACL		—	9-52	—
Display ACL Data	n/a	—	9-54	—
Delete an ACL	n/a	—	9-53	—
Configure an ACL from a TFTP Server	n/a	—	9-63	—
Enable ACL Logging	n/a	—	9-68	—
Show ACL Resources	n/a	—	9-19	—
Access-List Resources Help	n/a	—	9-18	—

ACL Applications

ACLs can filter traffic from a host, a group of hosts, or from entire subnets. Where it is necessary to apply ACLs to filter traffic from outside a network or subnet, applying ACLs at the edge of the network or subnet removes unwanted traffic as soon as possible, and thus helps to improve system performance. ACLs filter inbound traffic only and can rapidly consume switch resources. For these reasons, the best places to apply ACLs are on “edge” ports where ACLs are likely to be less complex and resource-intensive.

Optional Network Management Applications

ACLs through a RADIUS server can also be augmented using the Identity-Driven Management (IDM) application available for use with PCM. However, the features described in this chapter can be used without PCM or IDM support, if desired.

For information on configuring client authentication on the switch, refer to chapter 5, “RADIUS Authentication and Accounting”.

Optional PCM and IDM Applications

ProCurve Manager is a Windows-based network management solution for all manageable ProCurve devices. It provides network mapping and polling capabilities, device auto-discovery and topology, tools for device configuration and management, monitoring network traffic, and alerts and troubleshooting information for ProCurve networks.

ProCurve Identity Driven Manager (IDM) is an add-on module to the ProCurve Manager plus (PCM+) application. IDM extends the functionality of PCM+ to include authorization control features for edge devices in networks using RADIUS servers and Web-Authentication, MAC-Authentication, or 802.1X security protocols.

For more information, including electronic copies of the PCM and IDM manuals, visit the ProCurve Web site at **www.procurve.com**. (The PCM and IDM documentation is available under **Network Management** on the **Product manuals page** of the **Technical Support** area.)

General Application Options

Layer 3 IP filtering with Access Control Lists (ACLs) enables you to improve network performance and restrict network use by creating policies for:

- **Switch Management Access:** Permits or denies in-band management access. This includes preventing the use of certain TCP or UDP applications (such as Telnet, SSH, web browser, and SNMP) for transactions between specific source and destination IP addresses.
- **Application Access Security:** Eliminates inbound, unwanted IP, TCP, or UDP traffic by filtering packets where they enter the switch on specific physical ports or trunks.

This chapter describes how to configure, apply, and edit ACLs, and how to monitor the results of ACL actions.

Notes

ACLs can enhance network security by blocking selected IP traffic, and can serve as one aspect of maintaining network security. *However, because ACLs do not provide user or device authentication, or protection from malicious manipulation of data carried in IP packet transmissions, they should not be relied upon for a complete security solution.*

ACLs do not screen non-IP traffic such as AppleTalk and IPX.

For ACL filtering to take effect, configure an ACL and then assign it to the inbound traffic on a statically configured port or trunk.

Table 9-1. Comprehensive Command Summary

Action	Command	Page
Configuring Standard (Numbered) ACLs	ProCurve(config)# [no] access-list < 1-99 > < deny permit > < any host <src-ip-addr> src-ip-address/mask > ¹ [log] ²	9-39
Configuring Extended (Numbered) ACLs	ProCurve(config)# [no] access-list <100-199> < deny permit > ip < any host <src-ip-addr> src-ip-address/mask > ¹ [log] ²	9-44
	ProCurve(config)# [no] access-list < 100-199 > < deny permit > < tcp udp > < any host <src-ip-addr> src-ip-address/mask > ¹ [eq < src-port tcp/udp-id >] < any host <dest-ip-addr> dest-ip-address/mask > ¹ [eq < dest-port tcp/udp-id >] [log] ²	9-44
Configuring Standard (Named) ACLs	ProCurve(config)# [no] ip access-list standard < name-str 1-99 > ProCurve(config-std-nacl)# < deny permit > < any host <src-ip-addr> src-ip-address/mask > ¹ [log] ²	9-50 9-50
Configuring Extended (Named) ACLs	ProCurve(config)# [no] ip access-list extended < name-str 100-199 > ProCurve(config-std-nacl)# < deny permit > ip < any host <src-ip-addr> src-ip-address/mask > ¹ < any host <dest-ip-addr> dest-ip-address/mask > ¹ [log] ²	9-50 9-50
	ProCurve(config-std-nacl)# < deny permit > < tcp udp > < any host <src-ip-addr> src-ip-address/mask > ¹ [eq < tcp/udp-port-# well-known-port-name >] < any host <dest-ip-addr> dest-ip-address/mask > ¹ [eq < tcp/udp-port-# well-known-port-name >] [log] ²	9-50
Enabling or Disabling an ACL	ProCurve(config)# [no] interface < port-list > access-group < name-str 1-99 100-199 >	9-52
Deleting an ACL from the Switch	ProCurve(config)# no ip access-list < standard < name-str 1-99 > > ProCurve(config)# no ip access-list < extended < name-str 100-199 > >	9-53

Action	Command	Page
Displaying ACL Data	ProCurve(config)# show access-list	9-54
	ProCurve(config)# show access-list [<i>acl-name-string</i>]	
	ProCurve(config)# show access-list config	
	ProCurve(config)# show access-list ports < <i>port-list</i> >	
	ProCurve(config)# show access-list resources	
	ProCurve(config)# access-list resources help	
	ProCurve(config)# show config	
	ProCurve(config)# show running	

¹ The mask can be in either dotted-decimal notation (such as 0.0.15.255) or CIDR notation (such as /20).

² The [log] function applies only to “deny” ACLs, and generates a message only when there is a “deny” match.

Terminology

Access Control Entry (ACE): An ACE is a policy consisting of criteria and an action to take (permit or deny) on a packet if it meets the criteria. The elements composing the criteria include:

- Source IP address and mask (standard and extended ACLs)
- Destination IP address and mask (extended ACLs only)
- TCP or UDP application port numbers (optional, extended ACLs only)


Access Control List (ACL): A list (or set) consisting of one or more explicitly configured Access Control Entries (ACEs) and terminating with an implicit “deny” default which drops any packets that do not have a match with any explicit ACE in the named ACL. The two classes of ACLs are “standard” and “extended”. See “Standard ACL” and “Extended ACL”.

ACE: See “Access Control Entry”.

ACL: See “Access Control List”.

ACL ID: A number or alphanumeric string used to identify an ACL. A *standard* ACL ID can have either a number from 1 to 99 or an alphanumeric string. An *extended* ACL ID can have either a number from 100 to 199 or an alphanumeric string.

ACL Mask: Follows an IP address (source or destination) listed in an ACE to specify either a subnet or a group of devices. Defines which bits in a packet's corresponding IP addressing must exactly match the IP addressing in the ACE, and which bits need not match (wildcards). For example:



As shown above, zeros in an ACL mask specify an exact match requirement for IP addresses, and ones specify a wildcard. In this example, a matching IP address would be any address in the range 10.10.10.1-255. (See also “How an ACE Uses a Mask To Screen Packets for Matches” on page 9-25.)

DA: The acronym for *Destination IP Address*. In an IP packet, this is the destination IP address carried in the header, and identifies the destination intended by the packet's originator. In an extended ACE, this is the second of two IP addresses required by the ACE to determine whether there is a match between a packet and the ACE. See also “SA”.

Deny: An ACE configured with this action causes the switch to drop an inbound packet for which there is a match within an applicable ACL. As an option, you can configure the switch to generate a logging output to a Syslog server and a console session.)

Extended ACL: This type of Access Control List uses layer-3 IP criteria composed of source and destination IP addresses and (optionally) TCP or UDP port criteria to determine whether there is a match with an IP packet. You can apply an extended ACL to inbound traffic on a port or trunk, including any inbound traffic with a DA belonging to the switch itself. Extended ACLs require an identification number (ID) in the range of 100 - 199 or an alphanumeric name.

Implicit Deny: If the switch finds no matches between an inbound packet and the configured criteria in an applicable ACL, then the switch denies (drops) the packet with an implicit “deny IP any” operation. You can preempt the implicit “deny IP any” in a given ACL by configuring **permit any** (standard) or **permit IP any any** (extended) as the last explicit ACE in the ACL. Doing so permits an inbound packet that is not explicitly permitted or denied by other ACEs configured sequentially earlier in the ACL. Unless otherwise noted, “implicit deny IP any” refers to the “deny” action enforced by both standard and extended ACLs.

Inbound Traffic: For the purpose of defining where the switch applies ACLs to filter traffic, inbound traffic is any IP packet that:

- Enters the switch through a physical port.
- Has a destination IP address (DA) that meets either of these criteria:
 - The packet's DA is for an external device.
 - The packet's DA is for an IP address configured on the switch itself. (This increases your options for protecting the switch from unauthorized management access.)

Because ACLs are assigned to physical ports or port trunks, an ACL that filters inbound traffic on a particular port or trunk examines packets meeting the above criteria that enter the switch through that port or trunk.

Outbound Traffic: This is any traffic *leaving the switch* through a physical port or trunk. The switch does not apply ACLs to outbound traffic or internally where routed traffic moves between VLANs. That is, ACL operation is not affected by enabling or disabling routing on the switch. (Refer also to “ACL Inbound Application Points” on page 9-9.)

Permit: An ACE configured with this action allows a port or trunk to permit an inbound packet for which there is a match within an applicable ACL.

SA: The acronym for *Source IP Address*. In an IP packet, this is the source IP address carried in the IP header, and identifies the packet's sender. In an extended ACE, this is the first of two IP addresses used by the ACE to determine whether there is a match between a packet and the ACE. See also “DA”.

Standard ACL: This type of Access Control List uses layer-3 IP criteria of source IP address to determine whether there is a match with an inbound IP packet. You can apply a standard ACL to inbound traffic on a port or trunk, including any inbound traffic with a DA belonging to the switch itself. Standard ACLs require an identification number (ID) in the range of 1 - 99 or an alphanumeric name.

Wildcard: The part of a mask that indicates the bits in a packet's IP addressing that do not need to match the corresponding bits specified in an ACL. See also **ACL Mask** on page 9-7.

Overview

Types of IP ACLs

Standard ACL: Use a standard ACL when you need to permit or deny traffic based on source IP address. Standard ACLs are also useful when you need to quickly control a performance problem by limiting traffic from a subnet, group of devices, or a single device. (This can block all inbound IP traffic from the configured source, but does not block traffic from other sources within the network.) This ACL type uses a numeric ID of 1 through 99 or an alphanumeric ID string. You can specify a single host, a finite group of hosts, or any host.

Extended ACL: Use extended ACLs whenever simple IP source address restrictions do not provide the breadth of traffic selection criteria you want for a port or trunk. Extended ACLs allow use of the following criteria:

- Source and destination IP addresses
- TCP application criteria
- UDP application criteria

ACL Inbound Application Points

You can apply ACL filtering to IP traffic inbound on a physical port or static trunk with a destination (DA):

- On another device. (ACLs are not supported on dynamic LACP trunks.)
- On the switch itself. In figure 9-1, below, this would be any of the IP addresses shown in VLANs “A”, “B”, and “C” on the switch. (IP routing need not be enabled.)

The switch can apply ACL filtering to traffic *entering the switch* on ports and/or trunks configured to apply ACL filters. For example, in figure 9-1 you would assign an inbound ACL on port 1 to filter a packet from the workstation 10.28.10.5 to the server at 10.28.20.99. Note that all ACL filtering is performed on the inbound port or trunk. Routing may be enabled or disabled on the switch, and any permitted inbound traffic may have any valid destination.

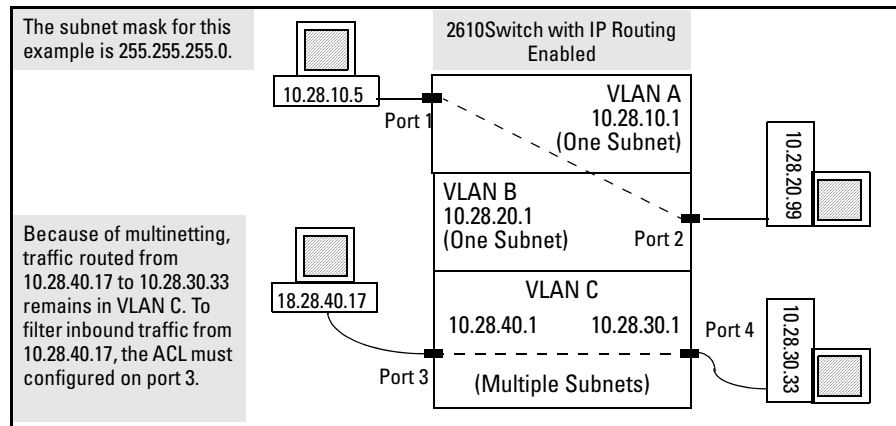


Figure 9-1. Example of Filter Applications

Features Common to All ACLs

- On any port or static trunk you can apply one ACL to inbound traffic.
- Any ACL can have multiple entries (ACEs).
- You can apply any one ACL to multiple ports and trunks.
- A source or destination IP address and a mask, together, can define a single host, a range of hosts, or all hosts.
- Before changing the content of an ACL assigned to one or more ports or trunks, you must first remove the ACL from those ports or trunks.
- Every standard ACL includes an implied **“deny any”** as the last entry, and every extended ACL includes an implied **“deny IP any any”** as the last entry. The switch applies this action to any packets that do not match other criteria in the ACL.
- In any ACL, you can apply an ACL log function to ACEs that have a “deny” action. The logging occurs when there is a match on a “deny” ACE. (The switch sends ACL logging output to Syslog and, optionally, to a console session.)
- Standard and Extended ACL features cannot be combined in one ACL.

You can configure ACLs using either the CLI or a text editor. The text-editor method is recommended when you plan to create or modify an ACL that has more entries than you can easily enter or edit using the CLI alone. Refer to “Editing ACLs and Creating an ACL Offline” on page 9-60.

General Steps for Planning and Configuring ACLs

1. Identify the traffic type to filter. Options include:
 - Any inbound IP traffic
 - Inbound TCP traffic only
 - Inbound UDP traffic only
2. The SA and/or the DA of inbound traffic you want to permit or deny.
3. Determine the best points at which to apply specific ACL controls. For example, you can improve network performance by filtering unwanted traffic at the edge of the network instead of in the core.
4. Design the ACLs for the selected control points. Where you are using explicit “deny” ACEs, you can optionally use the ACL logging feature to help verify that the switch is denying unwanted packets where intended. Remember that excessive ACL logging activity can degrade the switch's performance. (Refer to “Enable ACL “Deny” Logging” on page 9-67.)
5. Create the ACLs in the selected switches.
6. Assign the ACLs to filter the inbound traffic on ports and/or static trunk interfaces configured on the switch.
7. Test for desired results.

For more details on ACL planning considerations, refer to “Planning an ACL Application” on page 9-16.

Caution Regarding the Use of Source Routing

Source routing is enabled by default on the switch and can be used to override ACLs. For this reason, if you are using ACLs to enhance network security, the recommended action is to use the **no ip source-route** command to disable source routing on the switch. (If source routing is disabled in the running-config file, the **show running** command includes “**no ip source-route**” in the running-config file listing.)

ACL Operation

Introduction

An ACL is a list of one or more Access Control Entries (ACEs), where each ACE consists of a matching criteria and an action (permit or deny). An ACL applies only to the switch in which it is configured. ACLs operate on assigned ports and static trunks, and filter these traffic types:

- Traffic entering the switch. (Note that ACLs do not screen traffic at any internal point where traffic moves between VLANs or subnets within the switch; only on inbound ports and static trunks. Refer to “ACL Inbound Application Points” on page 9-9.)
- Switched or routed traffic entering the switch and having an IP address on the switch as the destination

You can apply one inbound ACL to each port and static trunk configured on the switch. The complete range of options includes:

- **No ACL** assigned. (In this case, all traffic entering the switch on the interface does so without any ACL filtering, which is the default.)
- **One ACL** assigned to filter the inbound traffic entering the switch on the interface.
- **Multiple Assignments for the same ACL.** (The switch allows one ACL assignment to an interface, but you can assign the same ACL to multiple interfaces.)

Note

On a given port or trunk, after you assign an ACL, the default action is to deny any traffic that is not specifically permitted by the ACL. (This applies only to the inbound traffic flow filtered by the ACL.)

The Packet-Filtering Process

Sequential Comparison and Action. When the switch uses an ACL to filter a packet, it sequentially compares each ACE's filtering criteria to the corresponding data in the packet until it finds a match.

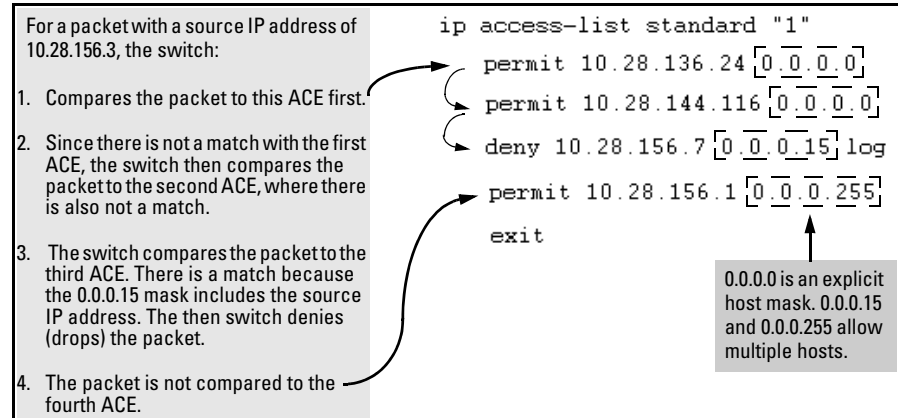


Figure 9-2. Example of Sequential Comparison

That is, the switch tries the first ACE in the list. If there is not a match, it tries the second ACE, and so on. When a match is found, the switch invokes the configured action for that entry (permit or drop the packet) and no further comparisons of the packet are made with the remaining ACEs in the ACL. This means that when the switch finds an ACE whose criteria matches a packet, it invokes the action configured for that ACE, and any remaining ACEs in the ACL are ignored. *Because of this sequential processing, successfully implementing an ACL depends in part on configuring ACEs in the correct order for the overall policy you want the ACL to enforce.*

Implicit Deny. If a packet does not have a match with the criteria in any of the ACEs in the ACL, the switch denies (drops) the packet. (This is termed *implicit deny*.) If you need to override the implicit deny so that any packet that does not have a match will be permitted, then you can enter **permit any** as the last ACE in the ACL. This directs the switch to permit (forward) any packets that do not have a match with any earlier ACE listed in the ACL, and prevents these packets from being filtered by the implicit deny.

Note on Implicit Deny

For ACLs configured to filter inbound packets, note that Implicit Deny filters *any packets, including those with a DA specifying the switch itself*. This operation helps to prevent management access from unauthorized IP sources.

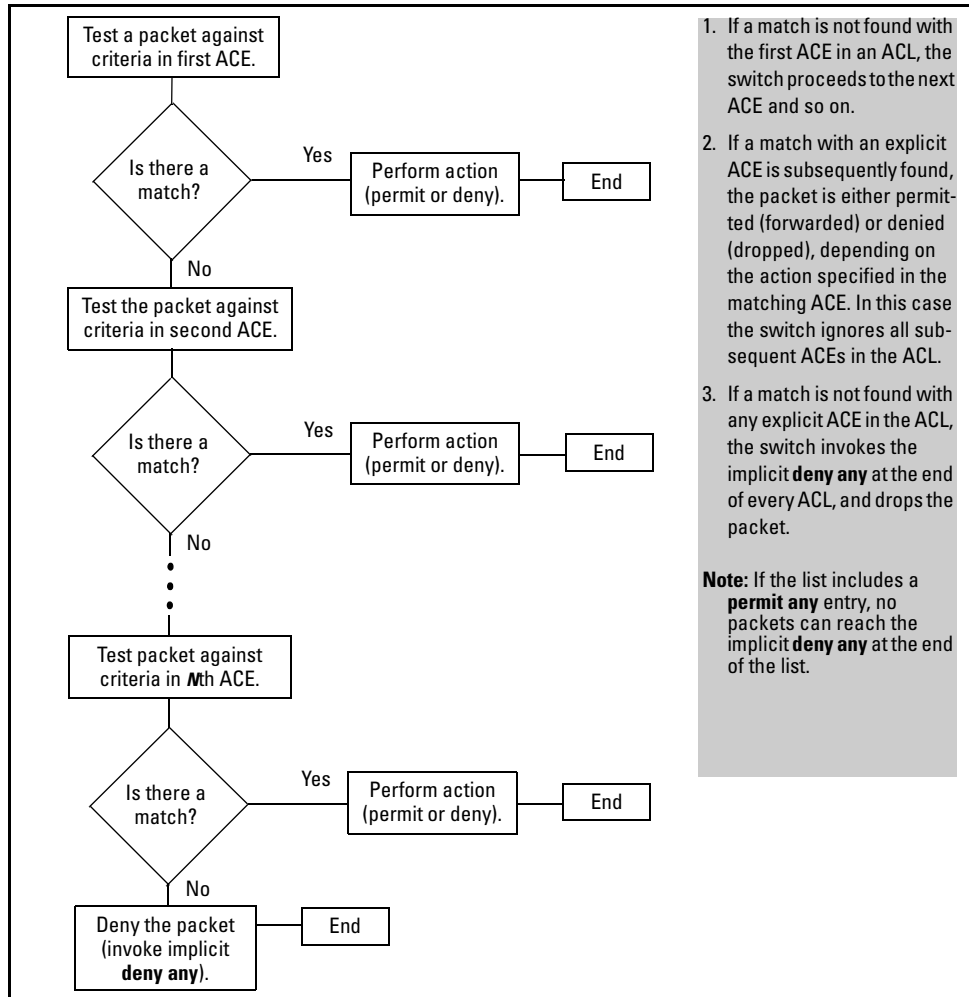


Figure 9-3. The Packet-Filtering Process in an ACL with N Entries (ACEs)

For example, suppose you want to configure an ACL on the switch (with an ID of “100”) to invoke these policies:

1. Permit all inbound traffic on port 12 sent from IP address 11.11.11.42.

2. Deny *only* the inbound Telnet traffic sent from IP address 11.11.11.101.
3. Permit *only* inbound Telnet traffic sent from IP address 11.11.11.33.
4. Deny *all other* inbound traffic on port 12.

The following ACL model, when assigned to inbound filtering on port 12, supports the above case:

```
ProCurve(config)# show access-list config

ip access-list extended "100"
  1 permit ip 11.11.11.42 0.0.0.0 0.0.0.0 255.255.255.255
  2 deny tcp 11.11.11.101 0.0.0.0 0.0.0.0 255.255.255.255 eq 23
  3 permit ip 11.11.11.101 0.0.0.0 0.0.0.0 255.255.255.255
  4 permit tcp 11.11.11.33 0.0.0.0 0.0.0.0 255.255.255.255 eq 23
  5 <implicit deny IP any >

ProCurve(config)# access-group 100 in
```

1. Permits IP traffic inbound from source address 11.11.11.42. Packets matching this criterion are permitted and will not be compared to any later ACE in the list. Packets not matching this criterion will be compared to the next entry in the list.	4. Permits Telnet traffic from source address 11.11.11.33. Packets matching this criterion are permitted and are not compared to any later criteria in the list. Packets not matching this criterion are compared to the next entry in the list.
2. Denies Telnet traffic from source address 11.11.11.101. Packets matching this criterion are dropped and are not compared to later criteria in the list. Packets not matching this criterion are compared to the next entry in the list.	5. This entry does not appear in an actual ACL, but is implicit as the last entry in every ACL. Any inbound packets on port 12 that do not match any of the criteria in the ACL's preceding entries will be denied (dropped).
3. Permits any IP traffic from source address 11.11.11.101. Any packets matching this criterion will be permitted and will not be compared to any later criteria in the list. Because this entry comes after the entry blocking Telnet traffic from this same address, there will not be any Telnet packets to compare with this entry; they have already been dropped as a result of matching the preceding entry.	

Figure 9-4. Example of How an ACL Filters Packets

It is important to remember that this ACL (and all ACLs) include an implicit **deny any**. That is, inbound IP packets (including switched packets having the switch as the destination IP address) that the ACL does not *explicitly* permit or deny will be *implicitly* denied, and therefore dropped. You can preempt the implicit deny by inserting a “permit IP any” at the end of an ACL, but this solution does not apply in the preceding example, where the intention is for the switch to allow only explicitly permitted packets inbound on port 12.

Overriding the Implicit “Deny Any”. If you want an ACL to permit any inbound packets that are not explicitly denied by other entries in the ACL, you can do so by configuring a **permit any** entry as the last entry in the ACL. Doing so permits any packet not explicitly denied by earlier entries. (On extended ACLs, you must configure **permit ip any any**.)

Planning an ACL Application

Before creating and implementing ACLs, you should understand the switch resources available to support ACL operation, define the policies you want your ACLs to enforce, and understand how your ACLs will impact your network users.

Switch Resource Usage

ACLs load resources in ways that require more careful attention to resource usage when planning a configuration using these features. Otherwise, there is an increased possibility of fully consuming some resources, which means that at some point the switch would not support further ACL configurations. This section describes resource planning for ACLs on your switch.

Prioritizing and Monitoring ACL and QoS, Feature Usage

If you want to configure ACLs on your switch, plan and implement your configuration in descending order of feature importance. This will help to ensure that the most important features are configured first. Also, if insufficient resources become a problem, this approach can help you recognize how to distribute the desired feature implementations across multiple switches to achieve your objectives.

ACL Resource Usage and Monitoring

ACL configurations use internal rules on a per-device basis. There are 128 rules available for configuring ACLs with the CLI and 128 rules available for configuring ACLs with IDM. You can apply a CLI ACL and an IDM ACL on the same port at the same time.

The switch uses resources required by the ACEs in an ACL when you apply the ACL to one or more port and/or static trunk interfaces.

Rule Usage

- There is only one implicit “deny any” entry per device for CLI ACLs, and one implicit “deny any” entry per device for IDM ACLs.
- The implicit “deny any” entry is created only the first time an ACL is applied to a port. After that the port-map is updated for that “deny any” entry to include or remove additional ports.
- Each ACE, including the implicit **deny any** ACE in a standard ACL, uses one rule.
- There is a separate rule for every ACE whether the ACE uses the same mask or a new mask.
- Two hardware rules are used for any “permit” ACE with TCP or UDP specified. One rule is for normal packets and one is for fragmented packets.

Table 9-2 on page 9-17 summarizes switch use of resources to support ACES.

Table 9-2. ACL Rule and Mask Resource Usage

ACE Type	Rule Usage
Standard ACLs	
Implicit deny any (automatically included in any standard ACL, but not displayed by show access-list < acl-# > command).	1
First ACE entered	1
Next ACE entered with same ACL mask	1
Next ACE entered with a different ACL mask	1
Closing ACL with a deny any or permit any ACE having the same ACL mask as the preceding ACE	1
Closing ACL with a deny any or permit any ACE having a different ACL mask than the preceding ACE	1
Extended ACLs	
Implicit deny ip any (automatically included in any standard ACL, but not displayed by show access-list < acl-# > command).	1
First ACE entered	1
Next ACE entered with same SA/DA ACL mask and same IP or TCP/UDP protocols specified	2
Next ACE entered with any of the following differences from preceding ACE in the list: <ul style="list-style-type: none"> – Different SA or DA ACL mask – Different protocol (IP as opposed to TCP/UDP) specified in either the SA or DA 	1
Closing an ACL with a deny ip any any or permit ip any any ACE preceded by an IP ACE with the same SA and DA ACL masks	1
Closing an ACL with a deny ip any any or permit ip any any ACE preceded by an IP ACE with different SA and/or DA ACL masks	1

The following two CLI commands are useful for planning and monitoring rule and mask usage in an ACL configuration.

Syntax: `access-list resources help`

Provides a quick reference on how ACLs use rule resources. Includes most of the information in table 9-2, plus an ACL usage summary.

Syntax: `show access-list resources`

Shows the number of rules used, maximum rules available, resources used and resources required for ACLs created with Identity Manager (IDM) and for ACLs created with the CLI.

Managing ACL Resource Consumption

As shown in table 9-2, changes in IP subnet masks or changes in IP or TCP/UDP applications among consecutive ACEs in an assigned ACL can rapidly consume resources. Adding a new ACE to an ACL consumes one rule. An extensive ACL configuration can fully subscribe the 128 rule resources available on the switch.

Oversubscribing Available Resources

If a given ACL requires more rule resources than are available, then the switch cannot apply the ACL to *any* of the interfaces specified for that ACL. In this case, the **access-group** command fails and the CLI displays the following:

- In the CLI:

Unable to apply access control list.

- In the Event Log (and in a Syslog server, if configured on the switch):

ACL: unable to apply ACL <acl-#> to port <port-#>, failed to add entry < # >

(Note that <port-#> is the first port in the assignment command that was unable to support the ACL.)

Troubleshooting a Shortage of Resources

Do the following to determine how to change resource usage to allow the ACL you want to configure:

1. Use the **show access-list resources** command
2. Use **show** commands to identify the currently configured ACL policies.

3. Determine which of the existing policies you can remove to free up rule resources for the ACL policy you want to implement. Depending on your network topology and configuration, you can free up rule resources by moving some policies to other devices. Another alternative is to inspect the switch's existing configuration for inefficient applications that could be removed or revised to achieve the desired policies with less resource usage. Table 9-2 on page 9-17 and the information displayed by the **access-list resources help** command, can help you to determine the resource usage of ACL policies.

Example of ACL Resource Usage

This example illustrates how to check for current rule availability, and then how to create and assign an ACL, and then to verify its effect on rule resources. (For more detailed information on configuring and applying ACLs, refer to the later sections of this chapter.)

Viewing the Current Rule Usage

The **show access-list resources** command displays current information about rules and resources.

```
ProCurve(config)# show access-list resources
ACL Resource Usage
```

Feature	Rules Used	Rules Maximum	Resources Used	Resources Required
cli-acl	15	128	1	1
idm-acl	0	128	0	2

Figure 9-5. Example of Rules Used and Resources Used and Required

Standard ACL Using a Subset of the Switch's Ports. Suppose that ports 1 - 4 belong to the following VLANs:

- VLAN 1: 10.10.10.1
- VLAN 2: 10.10.11.1
- VLAN 3: 10.10.12.1

(Assume that ports 1-4 are tagged members of VLAN 22, although tagged/untagged ports do not affect ACL operation because ACLs examine all inbound traffic, regardless of VLAN membership.)

The system administrator wants to:

- Permit inbound VLAN 1 traffic on all ports
- Permit inbound VLAN 2 traffic on ports 1 - 4 from hosts 10.10.10.1-30
- Deny inbound VLAN 2 traffic on ports 1 - 4 from hosts 10.10.10.31-255
- Permit inbound VLAN 3 traffic on all ports.

Because all ports in the example have the same inbound traffic requirements for ACL filtering, the system administrator needs to create only one ACL for application to all four ports.

- All inbound 10.10.10.x (VLAN 1) traffic is allowed on all ports.
- For the inbound 10.10.11.x (VLAN 2) traffic, the fourth octet of the ACL mask includes an overlap of permit and deny use on the “16” bit, which will require two different ACEs in the ACL. That is:
 - To deny hosts in the range of 31-255 in the fourth octet, it is necessary to use an ACE that specifies the leftmost four bits of the octet.
 - To permit hosts in the range of 1-30 in the fourth octet, it is necessary to use an ACE that specifies the rightmost five bits of the octet.

The overlap¹ can be illustrated as shown here:

Bit Values in the Fourth Octet	128	64	32	16	8	4	2	1
Bits Needed To Deny Hosts 31 - 255 (4th Octet Mask: 0.0.0.224)								
Bits Needed To Permit Hosts 1 - 30 (4th Octet Mask: 0.0.0.31)								
¹ For more on this topic, refer to “Rules for Defining a Match Between a Packet and an Access Control Entry (ACE)” on page 9-27, and “Using CIDR Notation To Enter the ACL Mask” on page 9-38.								

The overlap on the “16” bit means that it is necessary for the ACL to deny the host at 10.10.11.31 before permitting the hosts in the range of 10.10.10.1 - 30. The complete sequence is:

1. Permit all inbound traffic from 10.10.10.x.
2. Permit all inbound traffic from 10.10.12.x.
3. Deny the host at 10.10.11.31.

4. Permit the hosts in the range of 10.10.11.1 - 30.
5. Allow the implicit deny (automatically present in all ACLs) to deny all other traffic, which will automatically include the hosts in the range 10.10.10.32 - 255.

```
ProCurve(config)# access-list 1 permit 10.10.10.1/24
ProCurve(config)# access-list 1 permit 10.10.12.1/24
ProCurve(config)# access-list 1 deny host 10.10.11.31
ProCurve(config)# access-list 1 permit 10.10.11.1/27
ProCurve(config)# show access-list 1
```

Access Control Lists

Name: 1
Type: Standard
Applied: No

Every standard ACL has at least two ACEs; the first ACE that you configure, and the implicit **deny any** ACE that follows all other configured ACEs in the ACL.

ID	action		IP	Mask	Log
1	permit	std	10.10.10.1	0.0.0.255	
2	permit	std	10.10.12.1	0.0.0.255	
3	deny	std	10.10.11.31	0.0.0.0	
4	permit	std	10.10.11.1	0.0.0.31	

```
ProCurve(config)# interface 1-4 access-group 1 in
```

Figure 9-6. Example of Configuring an ACL

Traffic Management and Improved Network Performance

You can use ACLs to block unnecessary traffic caused by individual hosts, workgroups, or subnets, and to block user access to subnets, devices, and services. Answering the following questions can help you to design and properly position ACLs for optimum network usage.

- What are the logical points for minimizing unwanted traffic? In many cases it makes sense to block unwanted traffic from the core of your network by configuring ACLs to drop such traffic at or close to the edge of the network. (The earlier in the network path you block unwanted traffic, the greater the benefit for network performance.)
- What traffic should you explicitly block? Depending on your network size and the access requirements of individual hosts, this can involve creating a large number of ACEs in a given ACL (or a large number of ACLs), which increases the complexity of your solution and rapidly consumes the resources.
- What traffic can you implicitly block by taking advantage of the implicit **deny any** to deny traffic that you have not explicitly permitted? This can reduce the number of entries needed in an ACL and make more economical use of switch resources.
- What traffic should you permit? In some cases you will need to explicitly identify permitted traffic. In other cases, depending on your policies, you can insert a **permit any** (standard ACL) or **permit ip any any** (extended ACL) entry at the end of an ACL. This means that all IP traffic not specifically matched by earlier entries in the list will be permitted.

Security

ACLs can enhance security by blocking inbound IP traffic carrying an unauthorized source IP address (SA). This can include:

- Blocking access to or from subnets in your network
- Blocking access to or from the internet
- Blocking access to sensitive data storage or restricted equipment

- Preventing the use of specific TCP or UDP functions (such as Telnet, SSH, web browser) for unauthorized access

You can also enhance switch management security by using ACLs to block inbound IP traffic that has the switch itself as the destination address (DA).

Caution

ACLs can enhance network security by blocking selected IP traffic, and can serve as one aspect of maintaining network security. *However, because ACLs do not provide user or device authentication, or protection from malicious manipulation of data carried in IP packet transmissions, they should not be relied upon for a complete security solution.*

Note

ACLs do not screen non-IP traffic such as AppleTalk and IPX.

Guidelines for Planning the Structure of an ACL

The first step in planning a specific ACL is to determine where you will apply it. (Refer to “ACL Inbound Application Points” on page 9-9.) You must then determine the order in which you want the individual ACEs in the ACL to filter traffic. Some applications require high usage of the resources the switch uses to support ACLs. In these cases it is important to order the individual ACEs in a list to avoid unnecessarily using resources. For more on this topic, refer to “Planning an ACL Application” on page 9-16.

- The first match dictates the action on a packet. possible, subsequent matches are ignored.
- On any ACL, the switch implicitly denies packets that are not explicitly permitted or denied by the ACEs configured in the ACL. If you want the switch to forward a packet for which there is not a match in an ACL, add **permit any** as the last ACE in an ACL. This ensures that no packets reach the implicit **deny any** case.
- Generally, you should list ACEs from the most specific (individual hosts) to the most general (subnets or groups of subnets) unless doing so permits traffic that you want dropped. For example, an ACE allowing a small group of workstations to use a specialized printer should occur earlier in an ACL than an entry used to block widespread access to the same printer.

ACL Configuration and Operating Rules

- **Per-Interface ACL Limits.** At a minimum an ACL will have one explicit “deny” Access Control Entry. You can assign one ACL per interface, as follows:
 - Standard ACLs—Numeric range: 1 - 99
 - Extended ACLs—Numeric range: 100 - 199
 - Named (Extended or Standard) ACLs: Up to the maximum number of ports on the switch (minus any numeric ACL assignments)
- **Implicit “deny any”:** In any ACL, the switch automatically applies an implicit “deny IP any” that does not appear in **show** listings. This means that the ACL denies any packet it encounters that does not have a match with an entry in the ACL. Thus, if you want an ACL to permit any packets that you have not expressly denied, you must enter a **permit any** or **permit ip any any** as the last visible ACE in an ACL. Because, for a given packet the switch sequentially applies the ACEs in an ACL until it finds a match, any packet that reaches the **permit any** or **permit ip any any** entry will be permitted, and will not encounter the “deny ip any” ACE the switch automatically includes at the end of the ACL. For an example, refer to figure 9-4 on page 9-15.
- **Explicitly Permitting Any IP Traffic:** Entering a **permit any** or a **permit ip any any** ACE in an ACL permits all IP traffic not previously permitted or denied by that ACL.
- **Explicitly Denying Any IP Traffic:** Entering a **deny any** or a **deny ip any any** ACE in an ACL denies all IP traffic not previously permitted or denied by that ACL.
- **An ACL Assignment Is Exclusive:** The switch allows one ACL assignment on an interface. If a port or static trunk already has an ACL assigned, you cannot assign another ACL to the interface without first removing the currently assigned ACL.
- **Replacing One ACL with Another:** Where an ACL is already assigned to an interface, you must remove the current ACL assignment before assigning another ACL to that interface. If an assignment command fails because one or more interfaces specified in the command already have an ACL assignment, the switch generates this message in the CLI and in the Event Log:

<acl-list-#>: Unable to apply access control list.

- **ACLs Operate On Ports and Static Trunk Interfaces:** You can assign an ACL to any port and/or any statically configured trunk on the switch. ACLs do not operate with dynamic (LACP) trunks.
- **ACLs Screen Only the Traffic Entering the Switch on a Port or Static Trunk Interface:** On a given interface, ACLs can screen inbound traffic at the point where it enters the switch. ACLs do not screen traffic routed between VLANs within the switch, between subnets in a multinetted VLAN, or at the interface where the traffic exits from the switch. (See figure 9-1 on page 9-10.)
- **Before Modifying an Applied ACL, You Must First Remove It from All Assigned Interfaces:** An ACL cannot be changed while it is assigned to an interface.
- **Before Deleting an Applied ACL, You Must First Remove It from All Interfaces to Which It Is Assigned:** An assigned ACL cannot be deleted.
- **Port and Static Trunk Interfaces:**
 - Removing a port from an ACL-assigned trunk returns the port to its default settings.
 - To add a port to a trunk when an ACL is already assigned to the port, you must first remove the ACL assignment from the port.
 - Adding a new port to an ACL-assigned trunk automatically applies the ACL to the new port.

How an ACE Uses a Mask To Screen Packets for Matches

When the switch applies an ACL to inbound traffic on an interface, each ACE in the ACL uses an IP address and *ACL mask* to enforce a selection policy on the packets being screened. That is, the mask determines the range of IP addresses (SA only or SA/DA) that constitute a match between the policy and a packet being screened.

What Is the Difference Between Network (or Subnet) Masks and the Masks Used with ACLs?

In common IP addressing, a network (or subnet) mask defines which part of the IP address to use for the network number and which part to use for the hosts on the network. For example:

IP Address	Mask	Network Address	Host Address
18.38.252.195	255.255.255.0	first three octets	The fourth octet.
18.38.252.195	255.255.248.0	first two octets and the left-most five bits of the third octet	The right most three bits of the third octet and all bits in the fourth octet.

Thus, the bits set to 1 in a network mask define the part of an IP address to use for the network number, and the bits set to 0 in the mask define the part of the address to use for the host number.

In an ACL, IP addresses and masks provide the criteria for determining whether to deny or permit a packet, or to pass it to the next ACE in the list. If there is a match, the deny or permit action occurs. If there is not a match, the packet is compared with the next ACE in the ACL. Thus, where a standard network mask defines how to identify the network and host numbers in an IP address, the mask used with ACEs defines which bits in a packet's IP address must match the corresponding bits in the IP address listed in an ACE, and which bits can be *wildcards*.

Rules for Defining a Match Between a Packet and an Access Control Entry (ACE)

- For a given ACE, when the switch compares an IP address and corresponding mask in the ACE to an IP address carried in a packet:
 - **A mask-bit setting of 0 (“off”)** requires that the corresponding bit in the packet’s IP address and in the ACE’s IP address must be the same. That is, if a bit in the ACE’s IP address is set to 1 (“on”), the same bit in the packet’s IP address must also be 1.
 - **A mask-bit setting of 1 (“on”)** means the corresponding bit in the packet’s IP address and in the ACE’s IP address do not have to be the same. That is, if a bit in the ACE’s IP address is set to 1, the same bit in the packet’s IP address can be either 1 or 0 (“on” or “off”).

For an example, refer to “Example of How the Mask Bit Settings Define a Match” on page 9-29.

- In any ACE, a mask of all ones means *any* IP address is a match. Conversely, a mask of all zeros means the *only* match is an IP address identical to the host IP address specified in the ACL.
- Depending on your network, a single ACE that allows a match with more than one source or destination IP address may allow a match with multiple subnets. For example, in a network with a prefix of 31.30.240 and a subnet mask of 255.255.240.0 (the leftmost 20 bits), applying an ACL mask of 0.0.31.255 causes the subnet mask and the ACL mask to overlap one bit, which allows matches with hosts in two subnets: 31.30.224.0 and 31.30.240.0.

Bit Position in the Third Octet of Subnet Mask 255.255.240.0								
Bit Values	128	64	32	16	8	4	2	1
Subnet Mask Bits	1	1	1	1	n/a	n/a	n/a	n/a
Mask Bit Settings Affecting Subnet Addresses	0	0	0	1 or 0	n/a	n/a	n/a	n/a

This ACL supernetting technique can help to reduce the number of ACLs you need. You can apply it to a multinetted VLAN and to multiple VLANs. However, ensure that you exclude subnets that do not belong in the policy. If this creates a problem for your network, you can eliminate the unwanted match by making the ACEs in your ACL as specific as possible, and using multiple ACEs carefully ordered to eliminate unwanted matches.

- Every IP address and mask pair (source or destination) used in an ACE creates one of the following policies:

- **Any IP address fits the matching criteria.** In this case, the switch automatically enters the IP address and mask in the ACE. For example:

access-list 1 deny any

produces this policy in an ACL listing:

IP Address	Mask
0.0.0.0	255.255.255.255

This policy states that every bit in every octet of a packet's SA is a wildcard, which covers any IP address.

- **One IP address fits the matching criteria.** In this case, you provide the IP address and the switch provides the mask. For example:

access-list 1 permit host 18.28.100.15

produces this policy in an ACL listing:

IP Address	Mask
18.28.100.15	0.0.0.0

This policy states that every bit in every octet of a packet's SA must be the same as the corresponding bit in the SA defined in the ACE.

- **A group of IP addresses fits the matching criteria.** In this case you provide both the IP address and the mask. For example:

access-list 1 permit 18.28.32.1 0.0.0.31

IP Address	Mask
18.28.32.1	0.0.0.31

This policy states that:

- In the first three octets of a packet's SA, every bit must be set the same as the corresponding bit in the SA defined in the ACE.
- In the last octet of a packet's SA, the first three bits must be the same as in the ACE, but the last five bits are wildcards and can be any value.

- Unlike subnet masks, the wildcard bits in an ACL mask need not be contiguous. For example, 0.0.7.31 is a valid ACL mask. However, a subnet mask of 255.255.248.224 is not a valid subnet mask.

Example of How the Mask Bit Settings Define a Match . Assume an ACE where the second octet of the mask for an SA is 7 (the rightmost three bits are “on”, or “1”) and the second octet of the corresponding SA in the ACE is 31 (the rightmost five bits). In this case, a match occurs when the second octet of the SA in a packet being filtered has a value in the range of 24 to 31. Refer to table 9-3, below.

Table 9-3. Example of How the Mask Defines a Match

Location of Octet	Bit Position in the Octet							
	128	64	32	16	8	4	2	1
SA in ACE	0	0	0	1	1	1	1	1
Mask for SA	0	0	0	0	0	1	1	1
Corresponding Octet of a Packet's SA	0	0	0	1	1	0/1	0/1	0/1
The shaded area indicates bits in the packet that must exactly match the bits in the source IP in the ACE. Wherever the mask bits are ones (wildcards), the IP bits in the packet can be any value, and where the mask bits are zeros, the IP bits in the packet must be the same as the IP bits in the ACE. Note: This example covers only one octet of an IP address. An actual ACE applies this method to all four octets of an IP address.								

Example of Allowing Only One IP Address (“Host” Option). Suppose, for example, that you have configured the ACL in figure 9-7 to filter inbound packets on port 20. Because the mask is all zeros, the ACE policy dictates that a match occurs only when the source IP address on such packets is identical to the IP address configured in the ACE.

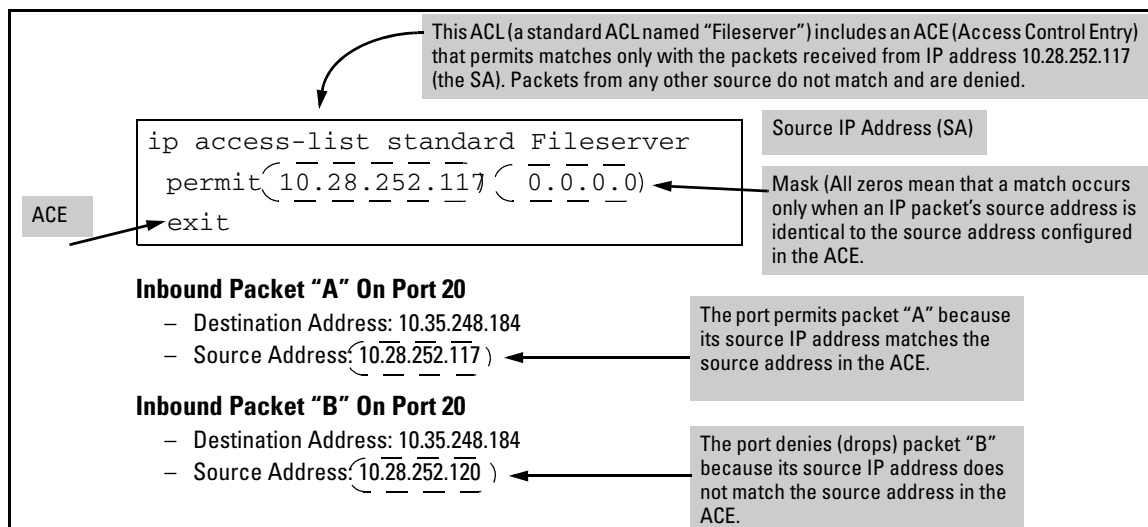


Figure 9-7. Example of an ACL with an Access Control Entry (ACE) that Allows Only One Source IP Address

Examples Allowing Multiple IP Addresses. Table 9-4 provides examples of how to apply masks to meet various filtering requirements.

Table 9-4. Example of Using an IP Address and Mask in an Access Control Entry

IP Address in the ACE	Mask	Policy for a Match Between a Packet and the ACE	Allowed IP Addresses
A: 10.38.252.195	0.0.0.255	Exact match in first three octets only.	10.38.252.< 0-255 > (See row A in table 9-5, below.)
B: 10.38.252.195	0.0.7.255	Exact match in the first two octets and the leftmost five bits (248) of the third octet.	10.38.< 248-255 >.< 0-255 > (In the third octet, only the rightmost three bits are wildcard bits. The leftmost five bits must be a match, and in the ACE, these bits are all set to 1. See row B in table 9-5, below.)
C: 10.38.252.195	0.0.0.0	Exact match in all octets.	10.38.252.195 (There are no wildcard bits in any of the octets. See row C in table 9-5, below.)
D: 10.38.252.195	0.15.255.255	Exact match in the first octet and the leftmost four bits of the second octet.	10.< 32-47 >.< 0-255 >.< 0-255 > (In the second octet, the rightmost four bits are wildcard bits. See row D in table 9-5, below.)

Table 9-5. Mask Effect on Selected Octets of the IP Addresses in Table 9-4

IP Addr	Octet	Mask	Octet Range	128	64	32	16	8	4	2	1
A	3	0 all bits	252	1	1	1	1	1	1	0	0
B	3	7 last 3 bits	248-255	1	1	1	1	1	0 or 1	0 or 1	0 or 1
C	4	0 all bits	195	1	1	0	0	0	0	1	1
D	2	15 last 4 bits	32-47	0	0	1	0	0 or 1	0 or 1	0 or 1	0 or 1

Shaded areas indicate bit settings that must be an exact match.

If there is a match between the policy in the ACE and the IP address in a packet, then the packet is either permitted or denied, according to how the ACE is configured. If there is not a match, the next ACE in the ACL is then applied to the packet. The same operation applies to a destination IP address (DA) used in an extended ACE. (Where an ACE includes both source and destination IP addresses, there is one IP-address/ACL-mask pair for the source address, and another IP-address/ACL-mask pair for the destination address. See “Configuring and Assigning an ACL” on page 9-32.)

CIDR Notation. For information on using CIDR notation to specify ACL masks, refer to “Using CIDR Notation To Enter the ACL Mask” on page 9-38.

Configuring and Assigning an ACL

ACL Feature	Page
Configuring and Assigning a Numbered, Standard ACL	9-39
Configuring and Assigning a Numbered, Extended ACL	9-44
Configuring a Named ACL	9-50
Enabling or Disabling ACL Filtering	9-52

Overview

General Steps for Implementing ACLs

1. Configure at least one ACL. This creates and stores the ACL in the switch configuration.
2. Assign an ACL. This applies the ACL to the inbound traffic on one or more designated interfaces.

Caution Regarding the Use of Source Routing

Source routing is enabled by default on the switch and can be used to override ACLs. For this reason, if you are using ACLs to enhance network security, the recommended action is to disable source routing on the switch. To do so, execute **no ip source-route**.

Types of ACLs

- **Standard ACL:** Uses only a packet's source IP address as a criterion for permitting or denying the packet. For a standard ACL ID, use either a unique numeric string in the range of 1-99 or a unique name string of up to 64 alphanumeric characters.
- **Extended ACL:** Offers the following criteria as options for permitting or denying a packet:
 - Source IP address
 - Destination IP address
 - TCP or UDP criteria

For an extended ACL ID, use either a unique number in the range of 100-199 or a unique name string of up to 64 alphanumeric characters.

You should carefully plan your ACL application before configuring specific ACLs. For more on this topic, refer to “Planning an ACL Application” on page 9-16.

ACL Configuration Structure

After you enter an ACL command, you may want to inspect the resulting configuration. This is especially true where you are entering multiple ACEs into an ACL. Also, it will be helpful to understand the configuration structure when using later sections in this chapter.

The basic ACL structure includes three elements:

- 1. ACL type and name: This identifies the ACL as **standard** or **extended** and shows the ACL name.
- 2. One or more deny/permit list entries (ACEs): One entry per line.

Element	Std	Ext	Notes
ID Range	1 - 99	100 - 199	You can also use an alphanumeric name of up to 64 characters, including spaces.
Minimum ACEs per ACL	1		
Maximum ACEs Per ACL	120		
Maximum ACEs per Switch	1024		
			In some cases, rule usage by ACLs may consume available resources to the point where this limit cannot be reached.

- 3. Implicit **deny any**: Where an ACL is in use, the switch denies any packets that do not have a match with the ACEs explicitly configured in the ACL. The implicit **deny any** does not appear in ACL configuration listings, but always functions when the switch uses an ACL to filter packets. (You cannot delete the implicit “deny any”, but you can supersede it with a “permit any” statement.)

Standard ACL Structure

Individual ACEs in a standard ACL include only a permit/deny “type” statement, the source IP addressing, and an optional **log** command (available with “deny” statements).

```
ip access-list < type > "< id-string >"
  permit host < source-ip-address >
  deny < source-ip-address > < acl-mask > [log]
  .
  .
  permit any
```

Figure 9-8. Example of the General Structure for a Standard ACL

For example, figure 9-9 shows how to interpret the entries in a standard ACL.

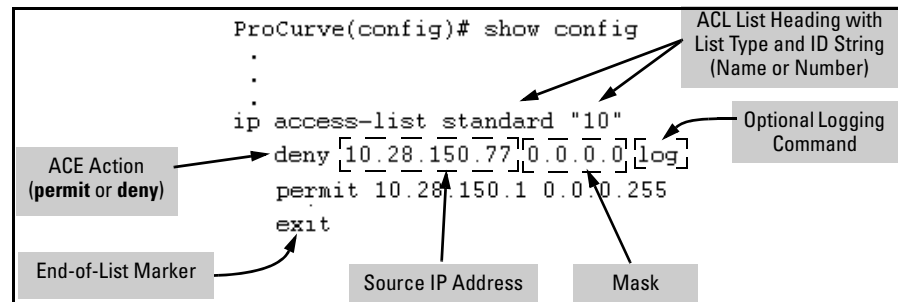


Figure 9-9. Example of a Displayed Standard ACL Configuration with Two ACEs

Extended ACL Configuration Structure

Individual ACEs in an extended ACL include:

- A permit/deny “type” statement
- Source IP addressing
- Optional TCP or UDP port type with optional source port ID and operator and/or optional destination port ID and operator
- Destination IP addressing
- Optional ACL **log** command (available for “Deny” ACLs only)

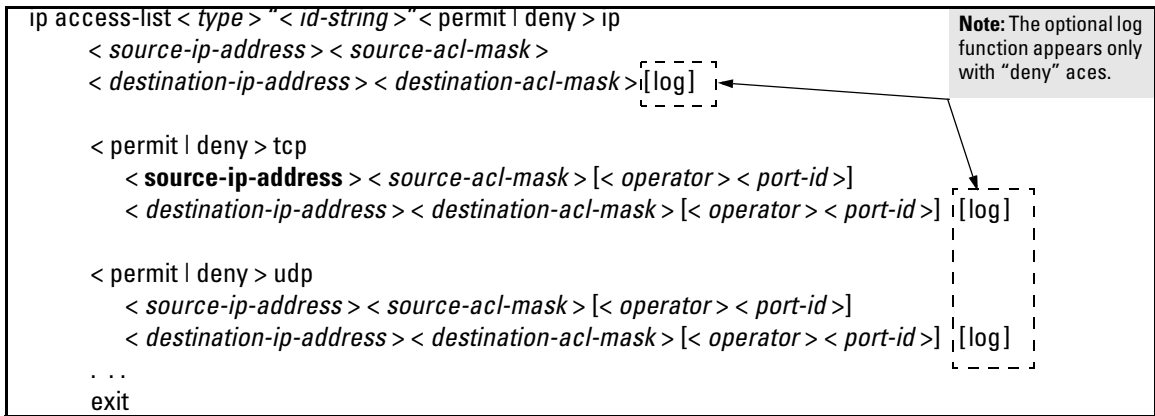


Figure 9-10. General Structure for an Extended ACL

For example, figure 9-11 shows how to interpret the entries in an extended ACL.

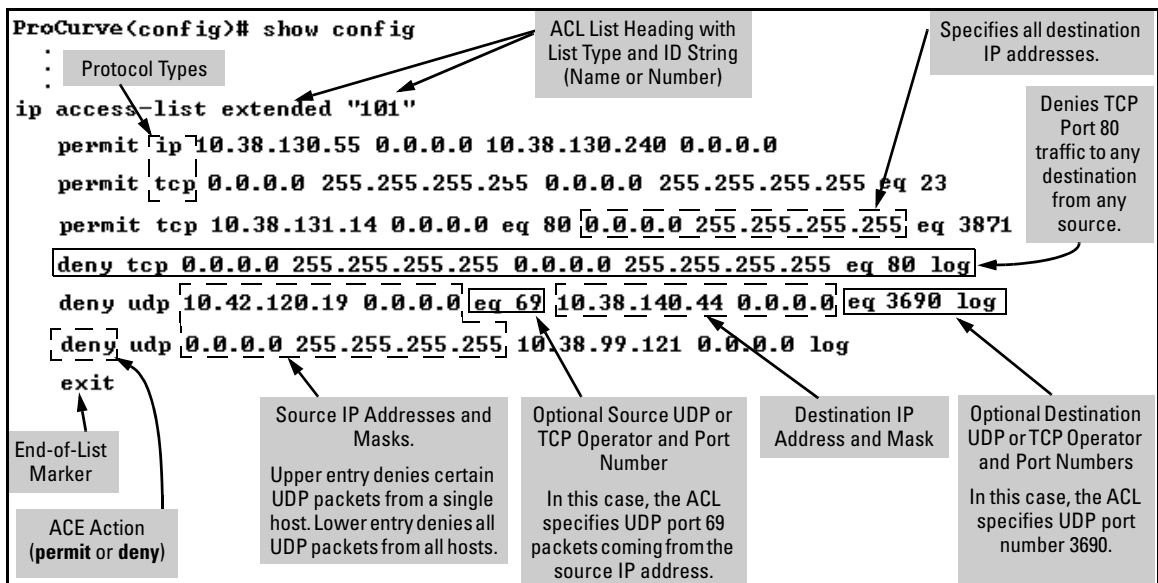


Figure 9-11. Example of a Displayed Extended ACL Configuration

ACL Configuration Factors

ACL Resource Consumption

Consumption of resources can be a significant factor in switches using extensive ACL applications. In this case, resource usage takes precedence over other factors when planning and configuring ACLs. For more information on this topic, refer to “Planning an ACL Application” on page 9-16.

The Sequence of Entries in an ACL Is Significant

When the switch uses an ACL to determine whether to permit or deny a packet on a particular interface, it compares the packet to the criteria specified in the individual Access Control Entries (ACEs) in the ACL, beginning with the first ACE in the list and proceeding sequentially until a match is found. When a match is found, the switch applies the indicated action (permit or deny) to the packet. This is significant because, once a match is found for a packet, subsequent ACEs in the same ACL will not be used for that packet, regardless of whether they match the packet.

For example, suppose that you have applied the ACL shown in figure 9-9 to inbound traffic on port 10:

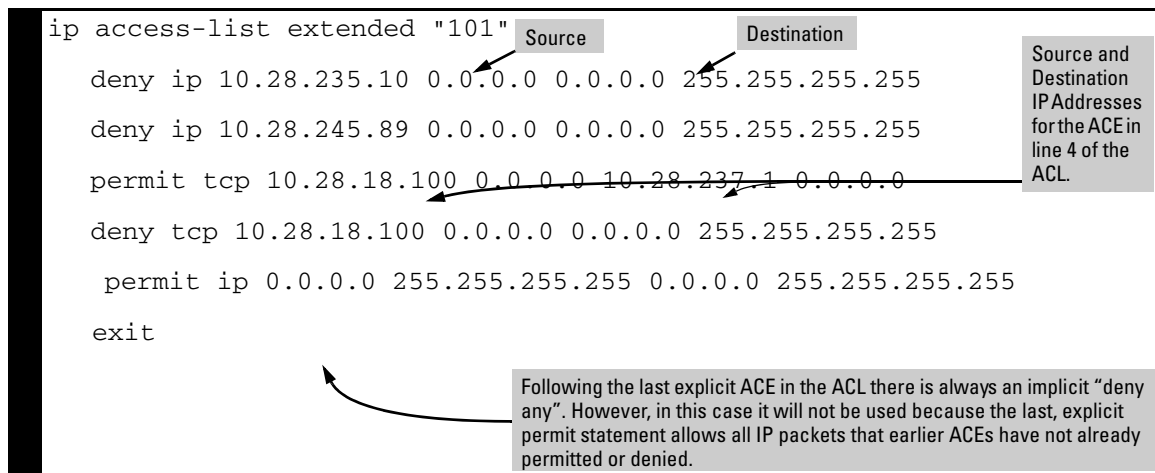


Figure 9-12. Example of an Extended ACL that Permits All Traffic Not Implicitly Denied

Table 9-6. Effect of the ACL in Figure 9-12 on Inbound Traffic on the Assigned Port

Line #	Action
1	Shows list type (extended) and ID (101).
2	A packet from IP source address 10.28.235.10 will be denied (dropped). This line filters out all packets received from 10.28.235.10. As a result, IP traffic from that device will not be routed or switched, and packets from that device will not be compared against any later entries in the list.
3	A packet from IP source 10.28.245.89 will be denied (dropped). This line filters out all packets received from 10.28.245.89. As the result, IP traffic from that device will not be routed or switched and packets from that device will not be compared against any later entries in the list.
4	A packet from TCP source address 10.28.18.100 with a destination address of 10.28.237.1 will be permitted (forwarded). Since no earlier lines in the list have filtered TCP packets from 10.28.18.100 and destined for 10.28.237.1, the switch will use this line to evaluate such packets. Any packets that meet this criteria will be forwarded. (Any packets that do not meet this TCP source-destination criteria are not affected by this line.)
5	A packet from TCP source address 10.28.18.100 to any destination address will be denied (dropped). Since, in this example, the intent is to block TCP traffic from 10.28.18.100 to any destination except the destination stated in line 4, this line must follow line 4. (If their relative positions were exchanged, all TCP traffic from 10.28.18.100 would be dropped, including the traffic for the 10.28.18.1 destination.)
6	Any packet from any IP source address to any destination address will be permitted (forwarded). The only traffic to reach this line will be IP packets not specifically permitted or denied in the earlier lines.
n/a	The “implicit deny any any” is a function automatically added as the last action in all ACLs. It denies (drops) any IP traffic from any source to any destination that has not found a match with earlier entries in the list. In this example, line 6 permits (forwards) any IP traffic not already permitted or denied by the earlier entries in the list, so there is no traffic remaining for action by the “implicit deny any any” function.
7	Indicates the end of the ACL.

In Any ACL, There Will Always Be a Match

As indicated in figure 9-12, the switch automatically uses an implicit “deny IP any” (Standard ACL) or “deny any” (Extended ACL) as the last ACE in any ACL. This means that if you configure the switch to use an ACL for filtering inbound traffic, any packets not specifically permitted or denied by the explicit entries you create will be denied by the implicit “deny” action. Note that if you want to preempt the implicit “deny” action, insert an explicit **permit any** or **permit ip any any** as the last line of the ACL.

A Configured ACL Has No Effect Until You Apply It to an Interface

The switch stores ACLs in the configuration file. Thus, until you actually assign an ACL to an interface, it is present in the configuration, but not used.

Using the CLI To Create an ACL

Command	Page
access-list (standard ACLs)	9-39
access-list (extended ACLs)	9-44
ip access-list (named ACLs)	9-50

You can use either the switch CLI or an offline text editor to create an ACL. This section describes the CLI method, which is recommended for creating short ACLs. (To use the offline method, refer to “Editing ACLs and Creating an ACL Offline” on page 9-60.)

General ACE Rules

These rules apply to all ACEs you create or edit using the CLI:

- ACEs are placed in an ACL according to the sequence in which you enter them (last entered, last listed).
- You can use the CLI to delete an ACE from anywhere in a given ACL by using the “no” form of the command to enter that ACE. However, when you use the CLI to add an ACE, the new entry is always placed *at the end of the ACL*.
- Duplicate ACEs are not allowed in an ACL, however the same ACE can be configured for multiple ACLs.

For more information, refer to “Editing ACLs and Creating an ACL Offline” on page 9-60.

Using CIDR Notation To Enter the ACL Mask

You can use CIDR (Classless Inter-Domain Routing) notation to enter ACL masks. The switch interprets the bits specified with CIDR notation as the IP address bits in an ACL and the corresponding IP address bits in a packet. The switch then converts the mask to inverse notation for ACL use.

Table 9-7. Examples of CIDR Notation for Masks

IP Address Used In an ACL with CIDR Notation	Resulting ACL Mask	Meaning
18.38.240.125/15	0.1.255.255	The leftmost 15 bits must match; the remaining bits are wildcards.
18.38.240.125/20	0.0.15.255	The leftmost 20 bits must match; the remaining bits are wildcards.
18.38.240.125/21	0.0.7.255	The leftmost 21 bits must match; the remaining bits are wildcards.
18.38.240.125/24	0.0.0.255	The leftmost 24 bits must match; the remaining bits are wildcards.
18.38.240.125/32	0.0.0.0	All bits must match.

Configuring and Assigning a Numbered, Standard ACL

To Configure:	Refer to:
Configuring Named ACLs	"Configuring a Named ACL" on page 9-50
Configuring Extended, Numbered ACLs	"Configuring and Assigning a Numbered, Extended ACL" on page 9-44

A standard ACL uses only source IP addresses in its ACEs. This type of ACE is useful when you need to:

- Permit or deny traffic based on source IP address only.
- Quickly control the IP traffic from a specific address, a group of addresses, or a subnet. This allows you to isolate traffic problems generated by a specific device, group of contiguous devices, or a subnet threatening to degrade network performance. This gives you an opportunity to troubleshoot without sacrificing performance for users outside of the problem area.

You can identify each standard ACL with a number in the range of 1 - 99, or an alphanumeric string of up to 64 characters. The CLI command process for using an alphanumeric string to name an ACL differs from the command process for a numeric name. For a description of how to name an ACL with an alphanumeric character string, refer to "Configuring a Named ACL" on page 9-50. To view the command differences, refer to table 9-1, "Comprehensive Command Summary" on page 9-5.

Note

For a summary of ACL commands, refer to table 9-1, “Comprehensive Command Summary”, on page 9-5.

Syntax: [no] access-list

Creates an ACE in the specified (1-99) access list and indicates the action (deny or permit) to take on a packet if there is a match between the packet and the criterion in the entry. If the ACL does not already exist, this command creates the specified ACL and its first ACE. To create a named ACL, refer to “Configuring a Named ACL” on page 9-50

< 1-99 >

Specifies the ACL ID number. The switch interprets an ACL with a value in this range as a standard ACL.

Note: To create an access list with an alphanumeric name (**name-str**) instead of a number, refer to “Configuring a Named ACL” on page 9-50.

< deny | permit >

Specifies whether to deny (drop) or permit (forward) a packet that matches the ACE criteria.

< any | host < src-ip-addr > | ip-addr / mask-length >

- **any**—Performs the specified action on any IP packet. Use this criterion to designate packets from any IP address.
- **host < host ip-address >**—Performs the specified action on any IP packet having the < host ip-address > as the source. Use this criterion to designate packets from a single IP address.

-
- **IP-addr / mask-length** — *Performs the specified action on any IP packet having a source address within the range defined by either*

< src-ip-addr / cidr-mask-bits >

or

< src-ip-addr < mask >>

Use this criterion to filter packets received from either a subnet or a group of contiguous IP addresses. The mask can be in either dotted-decimal format or CIDR format with the number of significant bits. Refer to “Using CIDR Notation To Enter the ACL Mask” on page 9-38.

The mask is applied to the IP address in the ACL to define which bits in a packet’s source IP address must exactly match the IP address configured in the ACL and which bits need not match. Note that specifying a group of contiguous IP addresses may require more than one ACE. For more on how masks operate in ACLs, refer to “How an ACE Uses a Mask To Screen Packets for Matches” on page 9-25.

[log]

Optionally generates an ACL log message if:

- *The action is **deny**.*
- *There is a match.*
- *ACL logging is enabled on the switch. (Refer to “Enable ACL “Deny” Logging” on page 9-67.)*

*(Use the debug command to direct ACL logging output to the current console session and/or to a Syslog server. Note that you must also use the **logging < ip-addr >** command to specify the IP addresses of Syslog servers to which you want log messages sent. See also “Enable ACL “Deny” Logging” on page 9-67.)*

Syntax: interface < port-list | trunk > access-group < ASCII-STR > in

Assigns an ACL, designated by an ACL ID (< ASCII-STR >), to an interface (list of one or more ports and/or one or more static trunks).

Access Control Lists (ACLs)

Configuring and Assigning an ACL

Example of a Standard ACL. Suppose you wanted to configure a standard ACL and assign it to filter inbound traffic on port 10 in a particular switch:

- The ID you selected for this ACL is “50”.
- You want the ACL to deny IP traffic from all hosts except these three:
 - 10.128.100.10
 - 10.128.100.27
 - 10.128.100.14

```
ProCurve(config)# access-list 50 permit host 10.128.100.10
ProCurve(config)# access-list 50 permit host 10.128.100.27
ProCurve(config)# access-list 50 permit host 10.128.80.14
ProCurve(config)# interface 10 access-group 50 in
ProCurve(config)# write mem
ProCurve(config)# show config
```

Startup configuration:

```
; J9085A Configuration Editor; Created on release #R.11.XX

hostname "ProCurve Switch 2610-24"
snmp-server contact "Allen Smith"
snmp-server location "Building P"
ip access-list standard "50"
  permit 10.128.100.10 0.0.0.0
  permit 10.128.100.27 0.0.0.0
  permit 10.128.80.14 0.0.0.0
  exit
interface 10
  access-group "50" in
exit
ip default-gateway 15.255.152.1
snmp-server community "public" Unrestricted
vlan 1
  name "DEFAULT_VLAN"
  untagged 1-28
  ip address dhcp-bootp
  exit
```

ProCurve(config)# show access-list resources

ACL Resource Usage

Feature	Rules Used	Rules Maximum	Resources Used	Resources Required
cli-acl	4	128	1	1
idm-acl	0	128	0	2

- Permits IP traffic from the indicated IP address. Since, for this example, ACL 50 is a new list, this command also creates the ACL.
- Permits IP traffic from the indicated IP address.
- The **deny any** that the switch implicitly includes in all standard ACLs denies IP packets from IP sources not included in the above three commands.

Show config lists any ACLs and ACL assignments configured in the startup-config.

ACL “50” is listed as assigned to filter inbound traffic on port 10.

show access-list resources shows the rule and resource usage.

Figure 9-13. Example of Configuring a Standard ACL To Permit Only Traffic from Specific IP Addresses

In a situation opposite to the above, suppose that you wanted to deny inbound IP traffic received on port 20 from 10.128.93.17 and 10.130.93.25, but permit all other IP traffic on this VLAN. The next ACL achieves this:

```
ProCurve Switch 2610-24(config)# access-list 60 deny host 10.128.93.17
ProCurve Switch 2610-24(config)# access-list 60 deny host 10.28.93.25
ProCurve Switch 2610-24(config)# access-list 60 permit any
ProCurve Switch 2610-24(config)# interface 20 access-group 60 in
ProCurve Switch 2610-24(config)# write mem
ProCurve Switch 2610-24(config)# show config
```

Startup configuration:

```
; J9085A Configuration Editor; Created on release #R.11.XX

hostname "ProCurve Switch 2610-24"
snmp-server contact "Allen Smith"
snmp-server location "Building P"
ip access-list standard "50"
| permit 10.128.100.10 0.0.0.0
| permit 10.128.100.27 0.0.0.0
| permit 10.128.80.14 0.0.0.0
| exit
ip access-list standard "60"
| deny 10.128.93.17 0.0.0.0
| deny 10.28.93.25 0.0.0.0
| permit 0.0.0.0 255.255.255.255
| exit
interface 10
| access-group "50" in
| exit
interface 20
| access-group "60" in
| exit
ip default-gateway 15.255.152.1
snmp-server community "public" Unrestricted
vlan 1
| name "DEFAULT_VLAN"
| untagged 1-28
| ip address dhcp-bootp
| exit
```

Show config lists any ACLs and ACL assignments configured in the startup-config.

ACL "50" from the preceding example.

ACL "60" is assigned to filter inbound traffic on port 20.

ACL "60" is listed in the switch configuration.

Denies IP traffic from the indicated IP address. Since, for this example, ACL 60 is a new list, this command also creates the ACL.

Denies IP traffic from the indicated IP address.

Permits IP traffic from all sources. (Traffic from the IP sources in the first two lines is already filtered and dropped.) The **deny any** with which the switch implicitly concludes all ACLs is preempted by this ACE (but is still present in the ACL).

Figure 9-14. Example of Configuring a Standard ACL To Deny Inbound Traffic from Specific IP Addresses

Configuring and Assigning a Numbered, Extended ACL

This section describes how to configure numbered, extended ACLs. To configure other ACL types, refer to the following table.

To Configure:	Refer To:
Standard, numbered ACLs	"Configuring and Assigning a Numbered, Standard ACL" on page 9-39
Named ACLs	"Configuring a Named ACL" on page 9-50

While standard ACLs use only source IP addresses for filtering criteria, extended ACLs allow multiple ACE criteria. This enables you to more closely define your IP packet-filtering criteria. These criteria include:

- Source and destination IP addresses (required), in one of the following options:
 - Specific host IP
 - Subnet or group of IP addresses
 - Any IP address
- IP protocol (IP, TCP, or UDP)
- Source TCP or UDP port (if the IP protocol is TCP or UDP)
- Destination TCP or UDP port (if the IP protocol is TCP or UDP)
- TCP or UDP **eq** operator (if the IP protocol is TCP or UDP)

You can configure extended ACLs with a numeric name in the range of 100 - 199. You can also configure extended ACLs with alphanumeric names. (Refer to "Configuring a Named ACL" on page 9-50.)

Note

For a summary of ACL commands, refer to table 9-1, "Comprehensive Command Summary", on page 9-5.

Syntax: [no] access-list

Creates an ACE in the specified (100-199) access list and:

- *Indicates the action (deny or permit) to take on a packet if there is a match between the packet and the criteria in the complete ACE.*
- *Specifies the packet protocol type (IP, TCP, or UDP).*
- *Specifies the source and destination addressing options described in the remainder of this section.*
- *Allows optional ACL logging where a packet has a match with a **deny** ACE.*

If the ACL does not already exist, this command creates the specified ACL and its first ACE. If the ACL already exists, this command adds a new, explicit ACE to the end of the ACL. For a match to occur, the packet must have the source and destination IP addressing criteria specified by this command, as well as any protocol-specific (TCP or UDP port number) criteria specified by the command. To create a named ACL, refer to “Configuring a Named ACL” on page 9-50.

< 100-199 >

Specifies the ACL ID number. The switch interprets an ACL with a value in this range as an extended ACL.

Note: *To create an access list with an alphanumeric name instead of a number, refer to “Configuring a Named ACL” on page 9-50.*

< deny | permit >

Specifies whether to deny (drop) or permit (forward) a packet that matches the ACE criteria.

< ip | tcp | udp >

Specifies the packet protocol type required for a match:

- **ip** — any IP packet
- **tcp** — only tcp packets
- **udp** — only udp packets

< any | host < src-ip-addr > | ip-addr/mask -length >

In an extended ACL, this parameter defines the source IP address (SA) that a packet must carry in order to have a match with the ACE.

- **any** — Specifies all inbound IP packets.
- **host < src-ip-addr >** — Specifies only inbound packets from a single IP address. Use this option when you want to match only the IP packets from one source IP address (device).
- **src-ip-addr/mask-length** — Performs the specified action on any IP packet having a source address within the range defined by either

< src-ip-addr / cidr-mask-bits >

or

< src-ip-addr < mask >>

Use this criterion to filter packets received from either a subnet or a group of IP addresses. The mask can be in either dotted-decimal format or CIDR format with the number of significant bits. Refer to “Using CIDR Notation To Enter the ACL Mask” on page 9-38.

The mask is applied to the IP address in the ACL to define which bits in a packet’s source IP address must exactly match the IP address configured in the ACL and which bits need not match. Note that specifying a group of contiguous IP addresses may require more than one ACE. For more on how masks operate in ACLs, refer to “How an ACE Uses a Mask To Screen Packets for Matches” on page 9-25.

[operator < src-port tcp/udp-id >]

*In an extended ACL where you have selected either **tcp** or **udp** as the packet protocol type (see above), you can optionally use a TCP or UDP source port number to further define the criteria for a match. To specify a TCP or UDP port number, (1) select the **eq** comparison operator and (2) enter the port number or a well-known port name.*

Comparison Operator:

- **eq** <tcp/udp-port-nbr> — “Equal To”; to have a match with the ACE entry, the TCP or UDP source port number in a packet must be equal to <tcp/udp-port-nbr>.

Port Number or Well-Known Port Name:

Use the TCP or UDP port number required by your application. The switch also accepts these well-known TCP or UDP port names as an alternative to their corresponding port numbers:

- **TCP:** bgp, dns, ftp, http, imap4, ldap, nntp, pop2, pop3, smtp, ssl, telnet
- **UDP:** bootpc, bootps, dns, ntp, radius, radius-old, rip, snmp, snmp-trap, tftp

To list the above names, press the **[Shift][?]** key combination after entering an operator. For a comprehensive listing of port numbers, visit www.iana.org/assignments/port-numbers.

< any | host < dest-ip-addr > | ip-addr/mask-length >

In an extended ACL, this parameter defines the destination IP address (DA) that a packet must carry in order to have a match with the ACE. The options are the same as shown for < **src-ip-addr** >.

[< dest-port tcp/udp-id >]

In an extended ACL, this parameter defines the TCP or UDP destination port number a packet must carry in order to have a match with the extended ACE. The options are the same as shown above on the preceding page for the source IP address.

[log]

Optional; generates an ACL log message if:

- The action is **deny**. (This option is not configurable for **Permit**.)
- There is a match.
- ACL logging is enabled on the switch. (Refer to “Enabling ACL Logging on the Switch” on page 9-68)

Syntax: interface < port-list > access-group < list-# | ascii-str > in

Assigns an ACL, designated by an ACL list number or ASCII string (alphanumeric list name), to an interface to filter inbound IP traffic on that interface. To configure named ACLs, refer to “Configuring a Named ACL” on page 9-50.

Example of an Extended ACL. Suppose that you want to implement these policies on ports 1, 2, and 3:

- A. Permit Telnet traffic from 10.10.10.44 inbound on port 1 to 10.10.20.78, deny all other inbound IP traffic from network 10.10.10.0 (VLAN 10) to 10.10.20.0 (VLAN 20), and permit all other IP traffic from any source to any destination. (See “A” in figure 9-15, below.)
- B. Permit FTP traffic from IP address 10.10.20.100 on port 2 to 10.10.30.55. Deny FTP traffic from other hosts on network 10.10.20.0 to any destination, but permit all other traffic.




Figure 9-15. Example of an Extended ACL


```
ProCurve(config)# access-list 110 permit tcp host 10.10.10.44 host 10.10.20.78
eq telnet
ProCurve(config)# access-list 110 deny ip 10.10.10.1/24 10.10.20.1/24
ProCurve(config)# access-list 110 permit ip any any
ProCurve(config)# interface 1 access-group 110 in
ProCurve(config)# access-list 120 permit tcp host 10.10.20.100 host 10.10.30.55
eq ftp
ProCurve(config)# access-list 120 deny tcp any any eq ftp
ProCurve(config)# access-list 120 permit ip any any
ProCurve config)# interface 2 access-group 120 in
ProCurve(config)# write mem
ProCurve(config)# show config
```

B (Refer to figure 9-15, above.)

write memory writes the configuration changes to the startup-config file.

Startup configuration:

```
; J9085A Configuration Editor; Created on release #R.11.XX
```

```
hostname "ProCurve Switch 2610-24"
snmp-server contact "Allen Smith"
snmp-server location "Building P"
ip access-list standard "50"
  permit 10.128.100.10 0.0.0.0
  permit 10.128.100.27 0.0.0.0
  permit 10.128.80.14 0.0.0.0
  exit
ip access-list standard "60"
  deny 10.128.93.17 0.0.0.0
  deny 10.28.93.25 0.0.0.0
  permit 0.0.0.0 255.255.255.255
  exit
```

```
ip access-list extended "110"
  permit tcp 10.10.10.44 0.0.0.0 10.10.20.78 0.0.0.0 eq 23
  deny ip 10.10.10.1 0.0.0.255 10.10.20.1 0.0.0.255
  permit ip 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255
  exit
ip access-list extended "120"
  permit tcp 10.10.20.100 0.0.0.0 10.10.30.55 0.0.0.0 eq 21
  deny tcp 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255 eq 21
  permit ip 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255
  exit
interface 1
  access-group "110" in
exit
interface 2
  access-group "120" in
exit
```

Access-List configuration in the switch's startup-config file.

Figure 9-16. Example of Configuration Commands for an Extended ACL

Configuring a Named ACL

You can use the “Named ACL” context to configure a standard or extended ACL with an alphanumeric name instead of a number. Note that the command structure for configuring a named ACL differs from that for a numbered ACL.

Syntax: ip access-list standard < name-str | 1-99 >
 < deny | permit >
 < any | host < src-ip-addr > | ip-addr / mask-length >
 [log]

ip access-list extended < name-str | 100-199 >
 < deny | permit > ip
 < any | host < src-ip-addr > | ip-addr / mask-length >
 < any | host < dest-ip-addr > | ip-addr / mask-length >
 [log]

ip access-list extended < name-string >
 < deny | permit > < tcp | udp >
 < any | host < src-ip-addr > | ip-addr / mask-length >
 [oper < src-port tcp/udp-id >]
 < any | host < dest-ip-addr > | ip-addr / mask-length >
 [oper < dest-port tcp/udp-id >]
 [log]

These commands create an ACE in the named ACL list and:

- *Indicate the action (deny or permit) to take on a packet if there is a match between a packet and the criteria in the complete ACE.*
- *Specify the packet protocol type (IP, TCP, or UDP) and (if TCP or UDP) the comparison operator.*
- *Specify the source and destination addressing options required for a match.*
- *Allow optional ACL logging where a packet has a match with a **deny** ACE. The **log** option does not appear when **permit** is the action.*

If the ACL does not already exist, these commands create the specified ACL and its first ACE. If the ACL already exists, these commands add a new, explicit ACE to the end of the ACL. For a match to occur, the packet must have the source and destination IP addressing criteria specified by this command, as well as any protocol-specific (TCP or UDP port number) criteria specified by the command.

< name-str | 1-99 | 100-199 >

Consists of an alphanumeric string of up to 64 case-sensitive characters. If you include a space in the string, you must also enclose the string with quotes. For example, "ACL # 1". You can also enter numbers in the ranges associated with standard (1-99) and extended (100-199) ACLs.

For explanations of the individual parameters in the preceding syntax statements, refer to the syntax descriptions under "Configuring and Assigning a Numbered, Standard ACL" on page 9-39 or "Configuring and Assigning a Numbered, Extended ACL" on page 9-44.

For example, figure 9-17 shows the commands for creating an ACL in the "Named ACL" context with these parameters:

ACL Name:	150
Action:	Deny
Protocol:	TCP
Source IP Address and Mask	10.10.20.100 0.0.0.0
Destination IP Address and Mask	10.10.10.1 0.0.0.255
Protocol Operator and Port Number at Destination	eq telnet

Access Control Lists (ACLs)

Configuring and Assigning an ACL

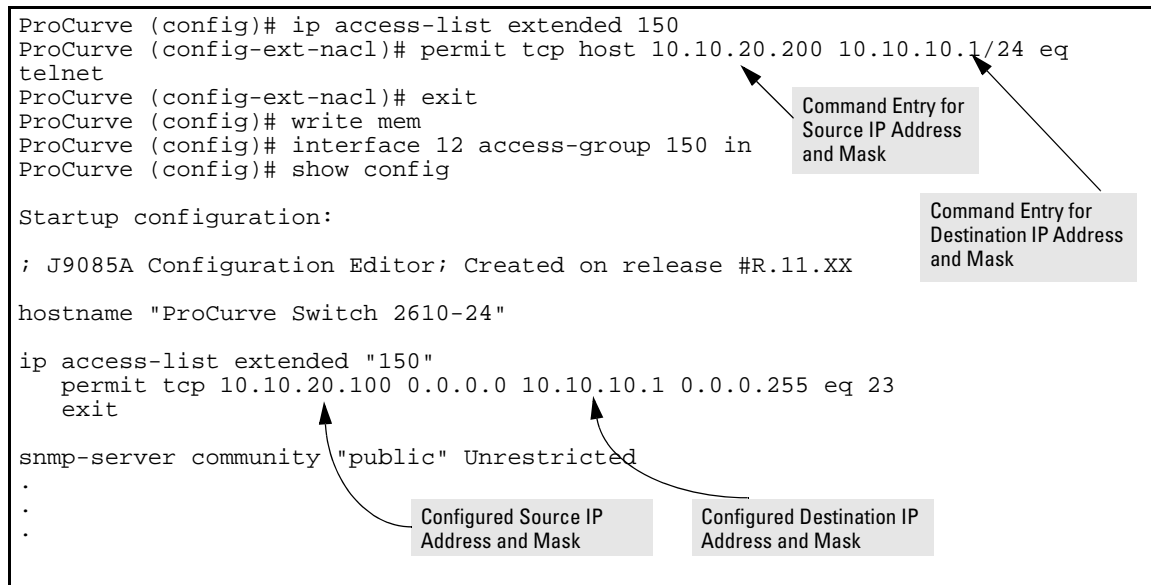


Figure 9-17. Using the “Named ACL” Context To Configure an ACL

Enabling or Disabling ACL Filtering on an Interface

You can configure one ACL to filter inbound traffic on multiple interfaces. For limits and operating rules, refer to “ACL Configuration and Operating Rules” on page 9-24.

Syntax: [no] interface < port-list > ip access-group < ascii-string > in
where: < ascii-string > = either a ACL name or an ACL ID number.

Assigns an ACL to a physical interface, which can be any combination of ports and/or trunks that do not already have an ACL assignment. You can use either the global configuration level or the interface context level to assign an ACL to an interface or remove an ACL from an interface.

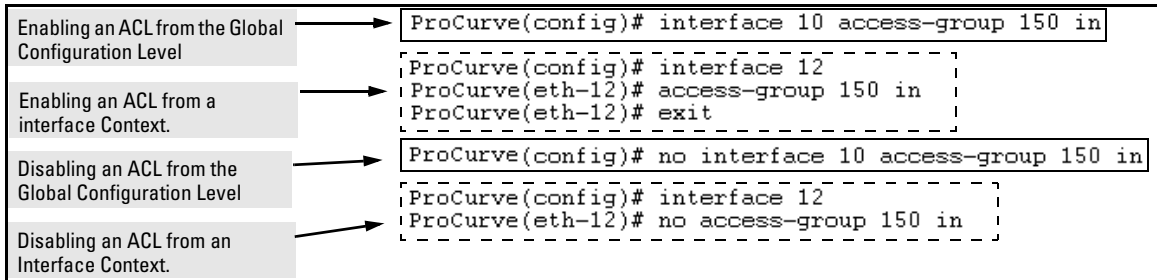


Figure 9-18. Methods for Enabling and Disabling ACLs

Deleting an ACL from the Switch

Syntax: no ip access-list standard < name-str | 1-99 >

no ip access-list extended < name-str | 100-199 >

Removes the specified ACL from the switch's running-config file.

Note: You cannot delete an ACL from the switch while the ACL is assigned to any interfaces. Thus, before deleting an ACL from the switch, remove all assignments of the ACL to specific interfaces. If you need to delete an ACL assignment, refer to “Enabling or Disabling ACL Filtering on an Interface” on page 9-52.

Displaying ACL Data

ACL Commands	Function	Page
show access-list	View a brief listing of all ACLs on the switch.	9-54
show access-list config	Display the ACL lists configured in the switch.	9-55
show access-list ports < all < interface >>	List the name and type of ACLs assigned to all ports on the switch or to a particular port or static trunk configured on the switch.	9-56
show access-list < acl-name-string >	Display detailed content information for a specific ACL.	9-57
show access-list resources	Displays the current rules and resources used.	9-59
show access-list radius	Displays ACLs applied via RADIUS	See chapter on RADIUS ACLs in this guide
show config	show config includes configured ACLs and assignments existing in the startup-config file.	
show running	show running includes configured ACLs and assignments existing in the running-config file.	

Display an ACL Summary

This command lists the configured ACLs, regardless of whether they are assigned to any interfaces.

Syntax: show access-list

List a summary table of the name, type, and application status of all ACLs configured on the switch.

For example:

```
ProCurve(config)# show access-list

Access Control Lists

Type  Appl  Name
----  -
ext   yes   110
ext   yes   120
ext   yes   150
std   yes   50
std   yes   60
```

Figure 9-19. Example of a Summary Table of Access Lists

Term	Meaning
Type	Shows whether the listed ACL is std (Standard; source-address only) or ext (Extended; protocol, source, and destination data).
Appl	Shows whether the listed ACL has been applied to an interface (yes/no).
Name	Shows the name or ID number assigned to each ACL configured in the switch.

Display the Content of All ACLs on the Switch

This command lists the configuration details for every ACL configured in the running-config file, regardless of whether you have assigned any to filter traffic on switch interfaces.

Syntax: show access-list config

List the configured syntax for all ACLs currently configured on the switch.

Note

Notice that you can use the output from this command for input to an offline text file in which you can edit, add, or delete ACL commands. Refer to “Editing ACLs and Creating an ACL Offline” on page 9-60.

This information also appears in the **show running** display. If you executed **write memory** after configuring an ACL, it appears in the **show config** display.

For example, with two ACLs configured in the switch, you will see results similar to the following:

```
ProCurve (config)# show access-list config

ip access-list standard "50"
  permit 10.128.100.10 0.0.0.0
  permit 10.128.100.27 0.0.0.0
  permit 10.128.80.14 0.0.0.0
  exit
ip access-list standard "60"
  deny 10.128.93.17 0.0.0.0
  deny 10.28.93.25 0.0.0.0
  permit 0.0.0.0 255.255.255.255
  exit
ip access-list extended "110"
  permit tcp 10.10.10.44 0.0.0.0 10.10.20.78 0.0.0.0 eq 23
  deny ip 10.10.10.1 0.0.0.255 10.10.20.1 0.0.0.255
  permit ip 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255
  exit
ip access-list extended "120"
  permit tcp 10.10.20.100 0.0.0.0 10.10.30.55 0.0.0.0 eq 21
  deny tcp 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255 eq 21
  permit ip 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255
  exit
ip access-list extended "150"
  permit tcp 10.10.20.100 0.0.0.0 10.10.10.1 0.0.0.255 eq 23
  permit tcp 10.10.20.200 0.0.0.0 10.10.10.1 0.0.0.255 eq 23
  exit
```

Figure 9-20. Example of an ACL Configured Syntax Listing

Display the ACL Assignments for an Interface

This command briefly lists the identification and type(s) of ACLs currently assigned to a particular interface (one or more ports and/or trunks) in the running-config file. (The switch allows up to one, inbound ACL assignment per interface.)

Syntax: show access-list ports < interface >

List the ACLs assigned to interfaces in the running config file.

Note

This information also appears in the **show running** display. If you executed **write memory** after configuring an ACL, it appears in the **show config** display.

For example, if you assigned a standard ACL with an ACL-ID of “1” to filter inbound traffic on port 10, you could quickly verify this assignment as follows:

```
ProCurve(config)# show access-list ports 7
[Access Lists for Port 7]
| Inbound   : 2 |
| Type      : Standard |
|_____|
ProCurve(config)# show access-list ports all
[Access Lists for Port 3]
| Inbound   : 1 |
| Type      : Standard |
|_____|
[Access Lists for Port 7]
| Inbound   : 2 |
| Type      : Standard |
|_____|
[Access Lists for Port Trk1]
| Inbound   : 2 |
| Type      : Standard |
|_____|
```

Indicates that a standard ACL with the ID of “2” is assigned to filter inbound traffic on port 7.

Indicates that a standard ACL with an ID of “1” is assigned to filter inbound traffic on port 3, and that another standard ACL with an ID of “2” is assigned to filter inbound traffic on port 7 and Trk1 (trunk 1).

Figure 9-21. Example of Listing the ACL Assignment for an Interface

Displaying the Content of a Specific ACL

This command displays a specific ACL configured in the running config file in an easy-to-read tabular format.

Note

This information also appears in the **show running** display. If you executed **write memory** after configuring an ACL, it appears in the **show config** display.

Syntax: show access-list < acl-name-string >

Display detailed information on the content of a specific ACL configured in the running-config file.

For example, suppose you configured the following two ACLs in the switch:

ACL ID	ACL Type	Desired Action
1	Standard	<ul style="list-style-type: none">• Deny IP traffic from 18.28.236.77 and 18.29.140.107.• Permit IP traffic from all other sources.
105	Extended	<ul style="list-style-type: none">• Permit any TCP traffic from 18.30.133.27 to any destination.• Deny any other IP traffic from 18.30.133.(1-255).• Permit all other IP traffic from any source to any destination.

Inspect the ACLs as follows:

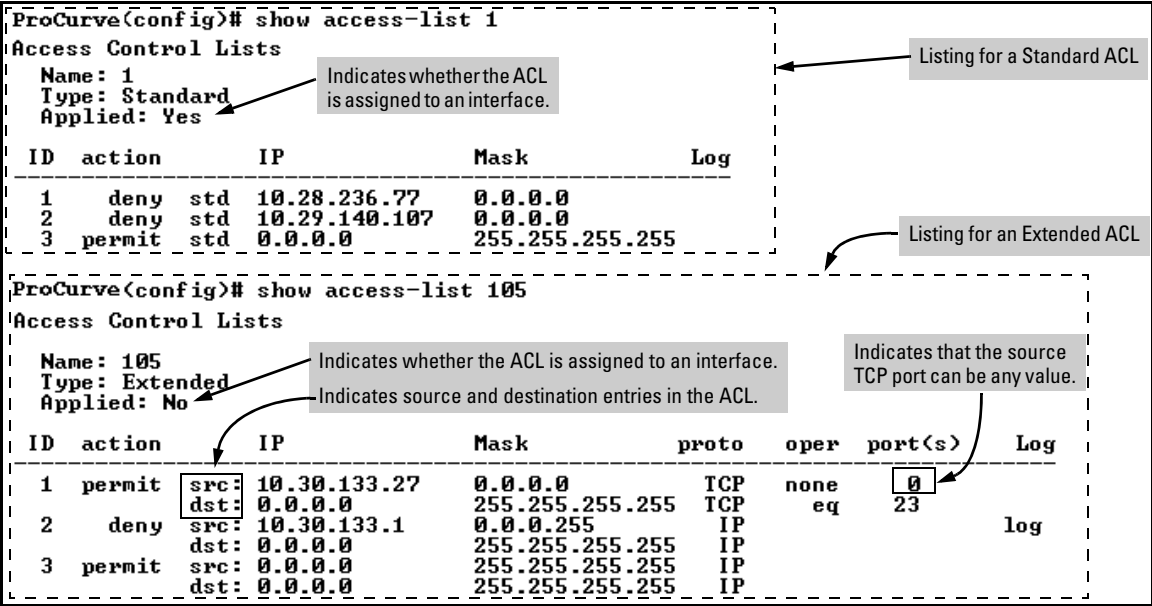


Figure 9-22. Examples of Listings Showing the Content of Standard and Extended ACLs

Table 9-8. Descriptions of Data Types Included in Show Access-List < interface > Output

Field	Description
Name	The ACL identifier. Can be a number from 1 to 199, or a name.
Type	Standard or Extended. The former uses only source IP addressing. The latter uses both source and destination IP addressing and also allows TCP or UDP port specifiers.
Applied	“Yes” means the ACL has been applied to an interface. “No” means the ACL exists in the switch configuration, but has not been applied to any interfaces, and is therefore not in use.
ID	The sequential number of the Access Control Entry (ACE) in the specified ACL.
action	Permit (forward) or deny (drop) a packet when it is compared to the criteria in the applicable ACE and found to match.
IP	In Standard ACLs: The source IP address to which the configured mask is applied to determine whether there is a match with a packet. In Extended ACLs: The source and destination IP addresses to which the corresponding configured masks are applied to determine whether there is a match with a packet.
Mask	The mask configured in an ACE and applied to the corresponding IP address in the ACE to determine whether a packet matches the filtering criteria.
proto	Used only in extended ACLs to specify the packet protocol type to filter. Must be either IP, TCP, or UDP.
oper	Used only in extended ACLs where a TCP or UDP port type and number have been entered. Specifies how to compare the corresponding TCP or UDP port number in a packet to the port number in the ACE.
port(s)	Used only in extended ACLs to show any TCP or UDP port number that has been entered in the ACE.
Log	Shows the status of logging for the entry (ACE). A blank space indicates ACL logging is not enabled for that ACE.

Displaying the Current ACL Resources

Assigning an ACL to one or more interfaces reduces the available resources for those interfaces. (An unassigned ACL does not affect the rule count.) This command displays the current rules and resources used on the switch. For more information on rule and mask usage, refer to “Planning an ACL Application” on page 9-16.

Syntax: show access-list resources

Displays the rules and resources that have been used on the switch. For more information, refer to “ACL Resource Usage and Monitoring” on page 9-16.

ProCurve(config)# show access-list resources				
ACL Resource Usage				
Feature	Rules Used	Rules Maximum	Resources Used	Resources Required
-----	-----	-----	-----	-----
cli-acl	15	128	1	1
idm-acl	0	128	0	2

Figure 9-23. Example of a Show Access-List Resources Command Output

Display All ACLs and Their Assignments in the Switch Startup-Config File and Running-Config File

The **show config** and **show running** commands include in their listings any configured ACLs and any ACL assignments to interfaces. Refer to figure 9-13 (page 9-42) and figure 9-14 (page 9-43) for examples. Remember that **show config** lists the startup-config file and **show running** lists the running-config file.

Editing ACLs and Creating an ACL Offline

Earlier sections of this chapter describe how to use the CLI to create an ACL. Beginning with “Using the CLI To Edit ACLs”, below, describes how to use the CLI to edit existing ACLs. However, you can also create or edit an ACL offline, then use a TFTP server to upload the ACL as a command file. The offline method (page 9-63) provides a useful alternative to using the CLI for creating or editing large ACLs.

Using the CLI To Edit ACLs

The switch applies individual ACEs in the order in which they occur in an ACL. You can use the CLI to delete individual ACEs from anywhere in an ACL and to append new ACEs to the end of an ACL. However, the CLI method does not allow you to insert a new ACE between two existing ACEs.

Note

Before editing an assigned ACL, you must use the **no interface < interface > access-group < acl-# > in** command to remove the ACL from all interfaces to which it is assigned.

Using the CLI To Edit a Short ACL. To insert a new ACE between existing ACEs in a short ACL, you may want to delete the ACL and then re-configure it by entering your updated list of ACEs in the correct order.

Using the CLI to Edit a Longer ACL. To insert a new ACE between existing ACEs in a longer ACL:

- a. Delete the first ACE that is out of sequence and all following ACEs through the end of the ACL.
- b. Re-Enter the desired ACEs in the correct sequence.

General Editing Rules

- You can delete any ACE from an ACL by repeating the ACE's entry command, preceded by the "no" statement. When you enter a new ACE, the switch inserts it as the last entry of the specified ACL.
- Deleting the last ACE from a *numeric* ACL, removes the ACL from the configuration. Deleting the last ACE from a *named* ACL leaves the ACL in memory. In this case, the ACL is "empty" and cannot perform any filtering tasks. (In any ACL the implicit "deny any" does not apply unless the ACL includes at least one explicit ACE.)
- When you create a new ACL, the switch inserts it as the last ACL in the startup-config file. (Executing **write memory** saves the running-config file to the startup-config file.)

Deleting Any ACE from an ACL

You can delete an ACE from an ACL by repeating the ACE's entry command, preceded by the "**no**" statement.

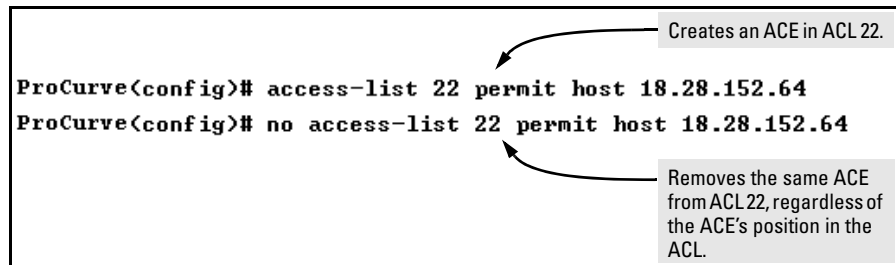
Syntax: no access-list < name-str /1-99> < permit | deny > < any | host | ip-addr/mask-length >

Deletes an ACE from a standard ACL. All variable parameters in the command must be an exact match with their counterparts in the ACE you want to delete.

```
no access-list < name-str / 100-199> < permit | deny> < ip | tcp | udp>  
< src-addr: any | host | ip-addr/mask-length> [operator < src-port-num>]  
< dest-addr: any | host | ip-addr-mask-length> [operator < dest-port-num>  
[log]
```

Deletes an ACE from an extended ACL. All variable parameters in the command must be an exact match with their counterparts in the ACE you want to delete.

For example, the first of the following two commands creates an ACE in ACL 22 and the second deletes the same ACE:



```
ProCurve(config)# access-list 22 permit host 18.28.152.64  
ProCurve(config)# no access-list 22 permit host 18.28.152.64
```

Creates an ACE in ACL 22.

Removes the same ACE from ACL 22, regardless of the ACE's position in the ACL.

Figure 9-24. Example of Deleting an ACE from a Standard ACL

Figure 9-25 shows an example of deleting an ACE from an extended ACL.

```
ProCurve(config)# show config
Startup configuration:
.
.
.
ip access-list extended "103"
deny tcp 0.0.0.0 255.255.255.255 10.10.10.2 0.0.0.0 eq 23 log
deny tcp 0.0.0.0 255.255.255.255 10.10.20.2 0.0.0.0 eq 23 log
permit ip 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255
exit
vlan 1
name "DEFAULT_VLAN"
untagged A1
.
.
.
ProCurve(config)# no access-list 103 deny tcp any host 10.10.20.2 eq 23 log
ProCurve(config)# write mem
ProCurve(config)# show config
Startup configuration:
.
.
.
ip access-list extended "103"
deny tcp 0.0.0.0 255.255.255.255 10.10.10.2 0.0.0.0 eq 23 log
permit ip 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255
exit
vlan 1
name "DEFAULT_VLAN"
untagged 1
```

ACL 103 Before Removing the Second "deny" ACE.

Use no access-list to remove this line from ACL 103.

ACL 103 After Removing the Second "deny" ACE.

Figure 9-25. Example of Deleting an ACE from an Extended ACL

Working Offline To Create or Edit an ACL

Note

When creating an ACL offline, ensure that there are sufficient rules available for the ACEs you plan to apply to the ACL. If you attempt to apply an ACL to multiple interfaces and the switch does not have sufficient resources to support the ACL, the command will fail for all specified interfaces. For more on ACL resources, refer to "Planning an ACL Application" on page 9-16.

For longer ACLs that would be difficult or time-consuming to accurately create or edit in the CLI, you can use the offline method:

1. Begin by doing one of the following:
 - To edit one or more existing ACLs, use **copy command-output tftp** to copy the current version of the ACL configuration to a file in your TFTP server. For example, to copy the ACL configuration to a file named **acl02.txt** in the TFTP directory on a server at 10.28.227.2:

```
ProCurve# copy command-output 'show access-list config' tftp 10.28.227.2 acl02.txt pc
```

- To create a new ACL, just open a text file in the appropriate directory on a TFTP server accessible to the switch.
- 2. Use the text editor to create or edit the ACL(s).
- 3. Use **copy tftp command-file** to download the file as a list of commands to the switch.

Creating an ACL Offline

Use a text editor that allows you to create an ASCII text file (.txt).

If you are replacing an ACL on the switch with a new ACL that uses the same number or name syntax, begin the command file with a “**no**” command to remove the earlier version of the ACL from the switch’s running-config file. Otherwise, the switch will append the new ACEs in the ACL you download to the existing ACL. For example, if you plan to use the **copy** command to *replace* ACL “103”, you would place this command at the beginning of the edited file:

```
no ip access-list extended 103
```




Figure 9-26. Example of an Offline ACL File Designed To Replace An Existing ACL

For example, suppose that you wanted to create an extended ACL to fulfill the following requirements (Assume a subnet mask of 255.255.255.0.):

- ID: 160
- Deny Telnet access to a server at 10.10.10.100 from these three IP addresses on port 2 (with ACL logging):
 - 10.10.20.17
 - 10.10.20.23
 - 10.10.20.40
- Allow any inbound access from all other addresses on port 2:

- Permit internet access to the following two IP addresses through port 24, but deny access to all other addresses through this port (without ACL logging).
 - 10.10.20.98
 - 10.10.20.21
 - Deny all traffic from port 3 to the server at 10.10.10.100 (without ACL logging).
 - Deny all traffic from port 5 to the server at 10.10.10.100 (without ACL logging), but allow any other traffic from port 5.
1. To create an ACL offline for the above requirements, you would create a **.txt** file with the content shown in figure 9-27.

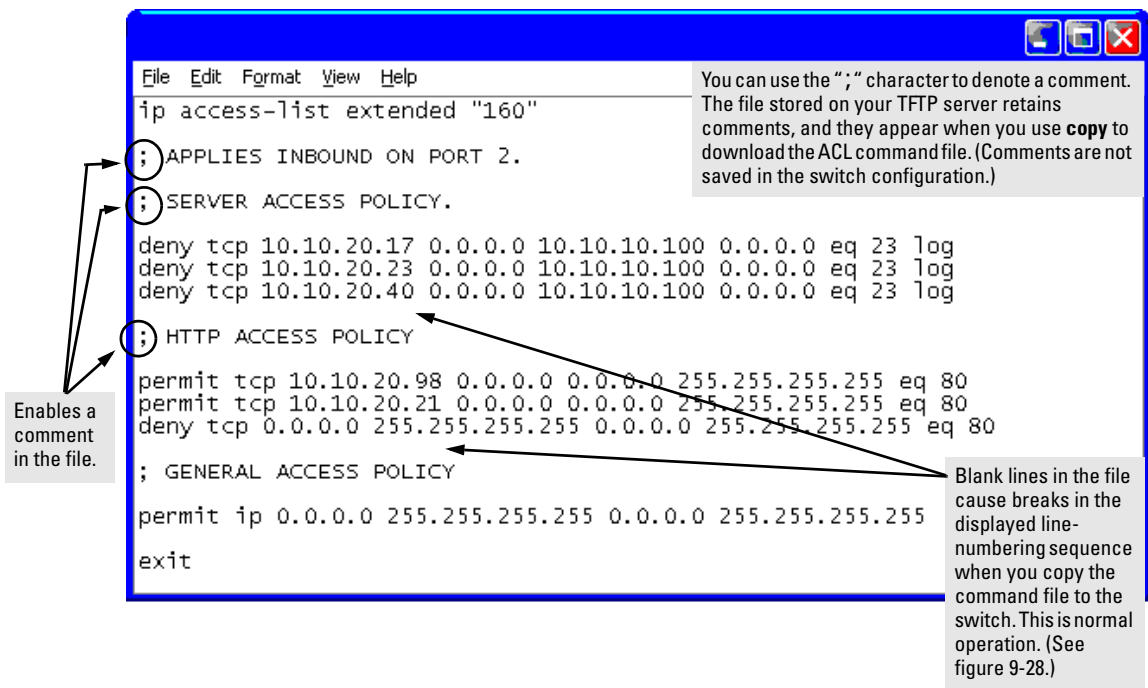


Figure 9-27. Example of a.txt File Designed for Creating an ACL

- After you copy the above .txt file to a TFTP server the switch can access, you would then execute the following command to download the file to the switch's startup-config file:

```
ProCurve(config)# copy tftp command-file 13.28.234.180 list-160.txt
Running configuration may change, do you want to continue [y/n]? y
 1. ip access-list extended "160"
 3. ; APPLIES INBOUND ON PORT 2.
 5. ; SERVER ACCESS POLICY.
 7. deny tcp 10.10.20.17 0.0.0.0 10.10.10.100 0.0.0.0 eq 23 log
 8. deny tcp 10.10.20.23 0.0.0.0 10.10.10.100 0.0.0.0 eq 23 log
 9. deny tcp 10.10.20.40 0.0.0.0 10.10.10.100 0.0.0.0 eq 23 log
11. ; HTTP ACCESS POLICY
13. permit tcp 10.10.20.98 0.0.0.0 0.0.0.0 255.255.255.255 eq 80
14. permit tcp 10.10.20.21 0.0.0.0 0.0.0.0 255.255.255.255 eq 80
15. deny tcp 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255 eq 80
17. ; GENERAL ACCESS POLICY
19. permit ip 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255
21. exit
```

Figure 9-28. Example of Using “copy tftp command-file” To Configure an ACL in the Switch

Note

If a transport error occurs, the switch does not execute the command and the ACL is not configured.

- Next, assign the new ACL to the intended interface which, in this example, is for port 2.

```
ProCurve(config)# interface 2 access-group 160 in
```

- Inspect the effect of the ACL on the switch's resources.

```
ProCurve(config)# show access-list resources
ACL Resource Usage
```

Feature	Rules Used	Rules Maximum	Resources Used	Resources Required
cli-acl	19	128	1	1
idm-acl	0	128	0	2

Figure 9-29. Inspection of Resource Usage After Assigning an ACL

- Inspect the new running configuration:

```
ProCurve(config)# show running
```

- If the configuration appears satisfactory, save it to the startup-config file:

```
ProCurve(config)# write memory
```

Enable ACL “Deny” Logging

ACL logging enables the switch to generate a message when IP traffic meets the criteria for a match with an ACE that results in an explicit “deny” action. You can use ACL logging to help:

- Test your network to ensure that your ACL configuration is detecting and denying the traffic you do not want forwarded
- Receive notification when the switch detects attempts to transmit traffic you have designed your ACLs to reject

The switch sends ACL messages to Syslog and optionally to the current console, Telnet, or SSH session. You can configure up to six Syslog server destinations.

Requirements for Using ACL Logging

- The switch configuration must include an ACL (1) assigned to an interface and (2) containing an ACE configured with the **deny** action and the **log** option.
- To screen routed packets with destination IP addresses outside of the switch, IP routing must be enabled.
- For ACL logging to a Syslog server, the server must be accessible to the switch and identified (with the **logging < ip-addr >** command) in the switch configuration.
- Debug must be enabled for ACLs and one or both of the following:
 - logging (for sending messages to Syslog)
 - Session (for sending messages to the current console interface)

ACL Logging Operation

When the switch detects a packet match with an ACE and the ACE includes both the **deny** action and the optional **log** parameter, an ACL log message is sent to the designated debug destination. The first time a packet matches an ACE with **deny** and **log** configured, the message is sent immediately to the destination and the switch starts a wait-period of approximately five minutes. (The exact duration of the period depends on how the packets are internally routed.) At the end of the collection period, the switch sends a single-line

summary of any additional “deny” matches for that ACE (and any other “deny” ACEs for which the switch detected a match). If no further log messages are generated in the wait-period, the switch suspends the timer and resets itself to send a message as soon as a new “deny” match occurs. The data in the message includes the information illustrated in figure 9-30.

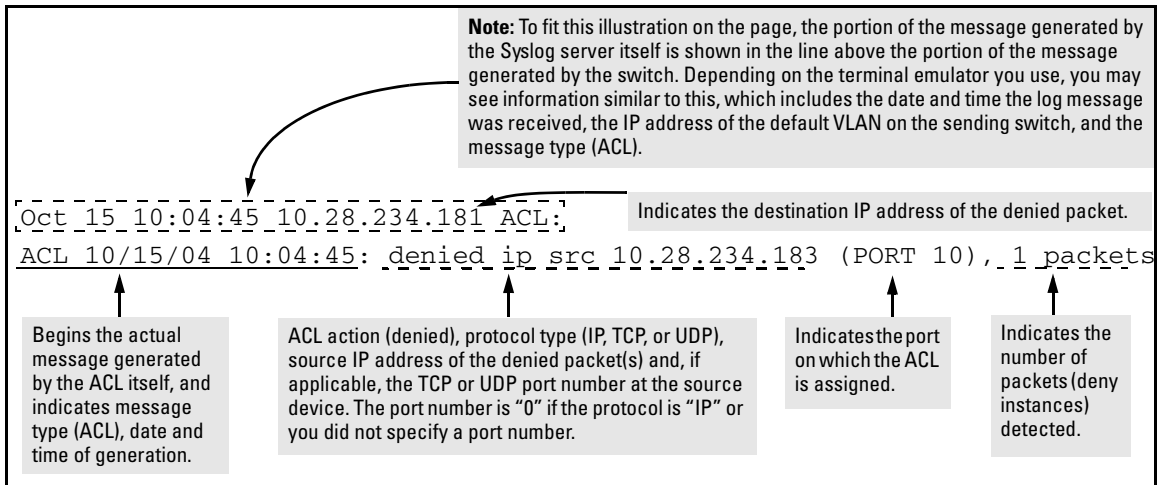


Figure 9-30. Example of the Content of an ACL-Generated Message

Enabling ACL Logging on the Switch

1. Use the debug command to:
 - a. Configure one or more log destinations.
 - b. If you are using a Syslog server, use the **logging** command to configure the server’s IP address. (You can configure up to six Syslog servers.)
 - c. Ensure that the switch can access any Syslog servers you specify.
2. Configure one or more ACLs with the deny action and the log option.

For example, suppose that you want to do the following:

- On port 10, configure an extended ACL with an ACL-ID of 143 to deny Telnet traffic from IP address 10.38.100.127.
- Configure the switch to send an ACL log message to the console and to a Syslog server at IP address 10.38.110.54 on port 11 if the switch detects a match denying Telnet access from 10.38.100.127.

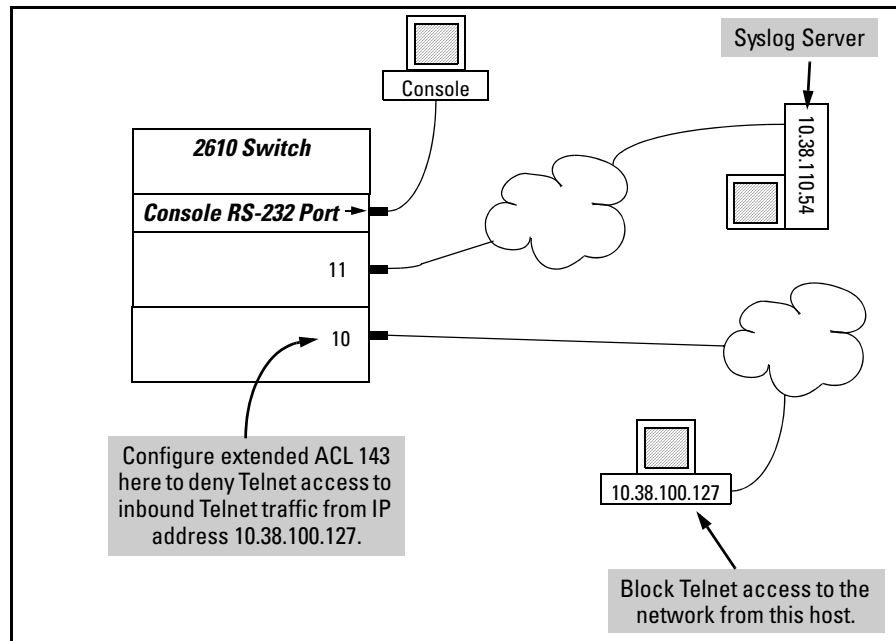


Figure 9-31. Example of an ACL Log Application

```
ProCurve(config)# access-list 143 deny tcp host 10.38.100.127 any eq telnet log
ProCurve(config)# access-list 143 permit ip any any
ProCurve(config)# interface 10 access-group 143 in
ProCurve(config)# logging 10.38.110.54
ProCurve(config)# debug acl
ProCurve(config)# debug destination logging
ProCurve(config)# debug destination session
ProCurve(config)# write memory
ProCurve(config)# show debug
Debug Logging
  Destination:
    Logging
      10.38.110.54
    Session
  Enabled debug types:
    event
    acl log
```

Figure 9-32. Commands for Applying an ACL with Logging to Figure 9-31

Operating Notes for ACL Logging

- The ACL logging feature generates a message only when packets are explicitly denied as the result of a match, and not when explicitly permitted or implicitly denied. To help test ACL logging, configure an ACL with an explicit **deny any** and **log** statements at the end of the list, and apply the ACL to an appropriate interface.
- Logging enables you to selectively test specific devices or groups. However, excessive logging can affect switch performance. For this reason, ProCurve recommends that you remove the logging option from ACEs for which you do not have a present need. Also, avoid configuring logging where it does not serve an immediate purpose. (Note that ACL logging is not designed to function as an accounting method.) See also “Apparent Failure To Log All “Deny” Matches” in the section titled “ACL Problems”, found in appendix C, “Troubleshooting” of the Management and Configuration Guide for your switch.
- When configuring logging, you can reduce excessive use by configuring the appropriate ACEs to match with specific hosts instead of entire subnets.

General ACL Operating Notes

ACLs do not provide DNS hostname support.

Protocol Support: ACL criteria includes IP, TCP, and UDP. ACLs do not use these protocols:

- TOS (Type-of-Service)
- Precedence
- MAC information
- QoS

ACLs do not affect switch serial port access.

ACLs filter both Layer 2 and Layer 3 on a port.

There is no performance degradation with ACLs enabled; traffic is at line rate.

When the ACL configuration includes TCP or UDP options, the switch operates in “strict” TCP and UDP mode for increased control. The switch compares all TCP and UDP packets against the ACLs.

Replacing or Adding To an Active ACL Policy. If you assign an ACL to an interface and subsequently want to add or replace ACEs in that ACL, you must first remove the ACL from all assigned interfaces.

Note

When an ACE becomes active, it screens the packets resulting from new traffic connections. It does not screen packets resulting from currently open traffic connections. If you invoke a new ACE to screen packets in a currently open traffic connection, you must force the connection to close before the ACE can begin screening packets from that source.

ACLs Do Not Filter Traffic Generated by the Switch. Because ACLs filter only inbound traffic at the inbound physical port, outbound traffic from any source is not filtered by any ACL(s) configured on the switch. Filtering of such traffic must be done at a downstream device.

`< acl-list-# >`: Unable to apply access control list.

The indicated ACL cannot be applied to an interface because an ACL is already assigned to the interface. The command fails for all included interfaces, including any that do not already have an ACL assigned.

Duplicate access control entry.

The switch detects an attempt to create a duplicate ACE in the same ACL.

Traffic/Security Filters

Contents

Overview	10-2
General Operation	10-2
Applying a Source Port Filter in a Multinetted VLAN	10-3
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	10-7
Filter Indexing	10-9
Editing a Source-Port Filter	10-9
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.

- With routing disabled on the switch (the default), source-port filtering can operate on traffic moving within the same VLAN.
- With routing enabled on the switch, source-port filtering can operate on traffic moving between VLANs as well as within the same VLAN. (If you configure multinetting within a VLAN and enable routing on the switch, you can use source-port filtering to filter traffic between subnets within the same VLAN.)
- 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 10-6.

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

Applying a Source Port Filter in a Multinetted VLAN

If you have multiple IP addresses configured on the same VLAN (multinetting), and routing is enabled on the switch, then a single port or trunk can be both the source and destination of packets moving between subnets in that same VLAN. In this case, you can prevent the traffic of one subnet from being routed to another subnet on the same port by configuring the port or trunk as both the source and destination for traffic to drop.

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 < source-port-number > [drop [forward] | forward
[drop]]

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 < 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 < 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:

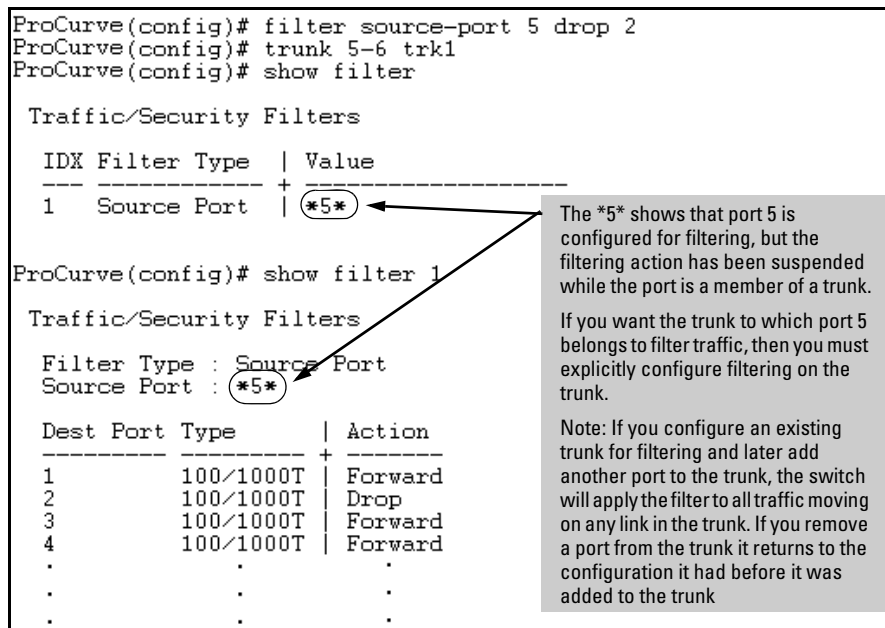


Figure 10-3. Example of Switch Response to Adding a Filtered Source Port to a Trunk

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

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.

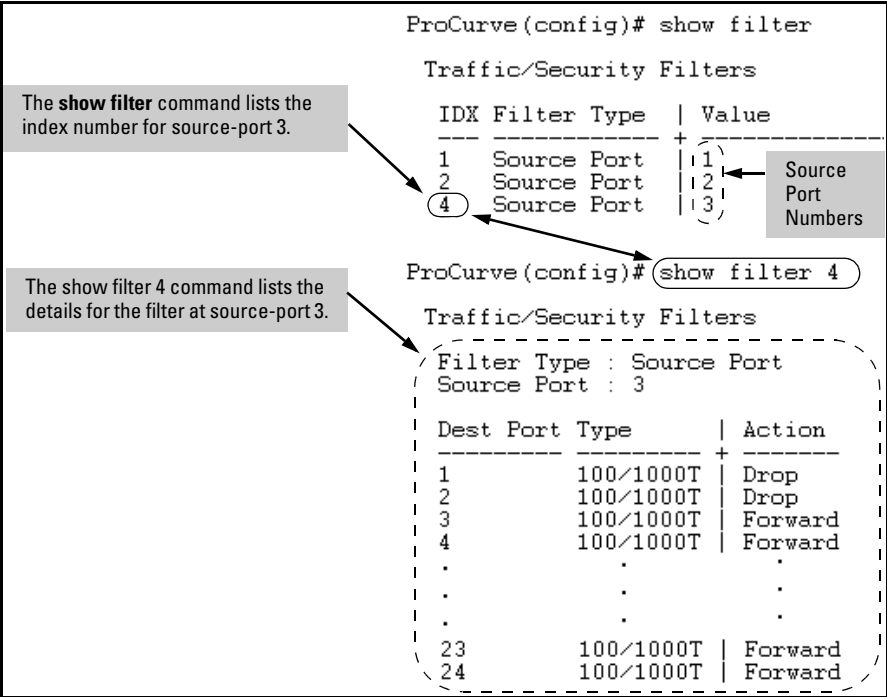


Figure 10-4. Example of Listing Filters and the Details of a Specific Filter

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.

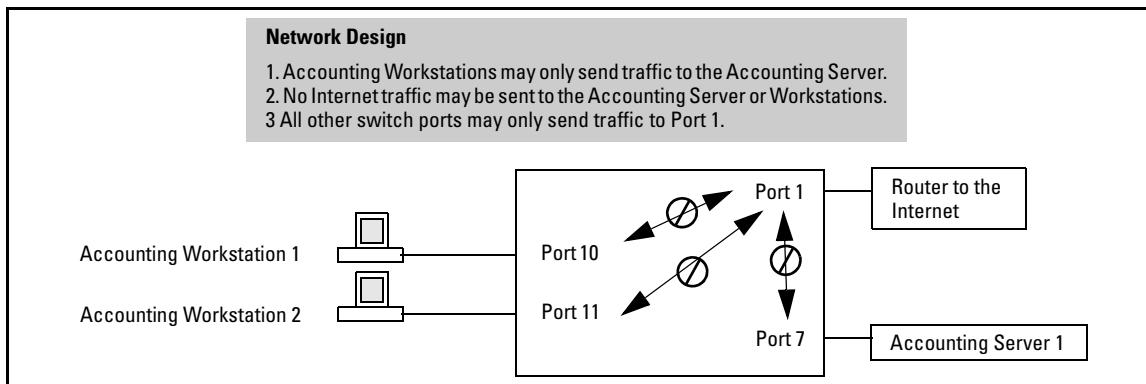


Figure 10-6. Network Configuration for Named Source-Port Filters Example

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.

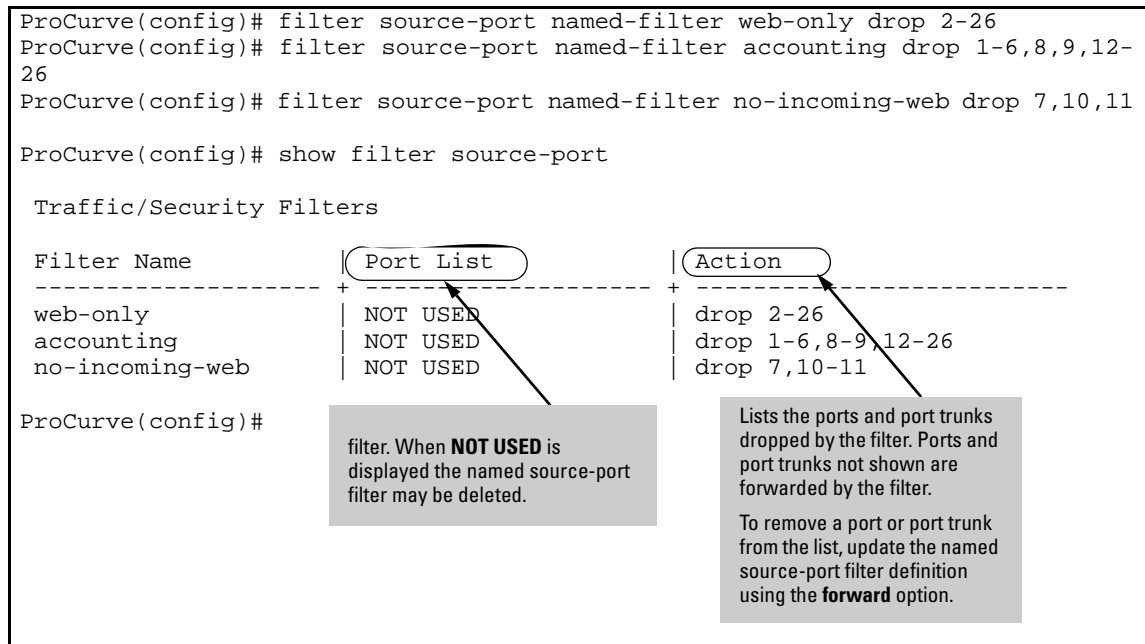


Figure 10-7. Example of filter source-port named-filter Command

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. Named source-port Filters Applied to Ports

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

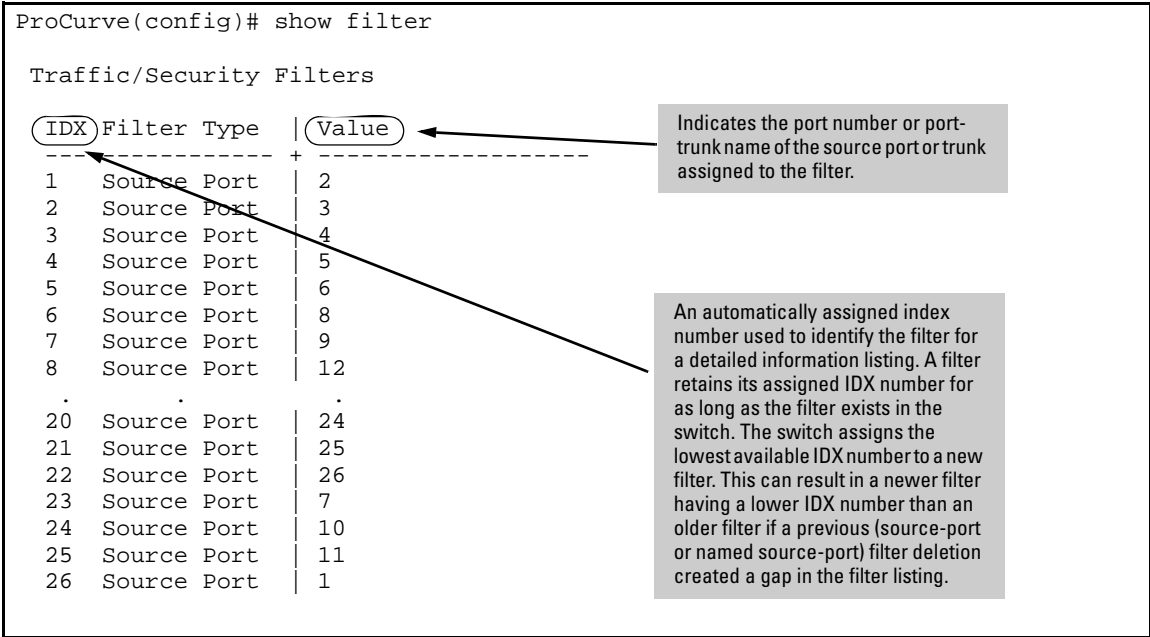


Figure 10-9. Show Command for Source Port Filters

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

Traffic/Security Filters

Filter Type : Source Port

Source Port : 5

Dest	Port	Type	Action
-----	-----	+	-----
1	10/100TX		Forward
2	10/100TX		Drop
3	10/100TX		Drop
4	10/100TX		Drop
5	10/100TX		Drop
6	10/100TX		Drop
7	10/100TX		Drop
8	10/100TX		Drop
9	10/100TX		Drop
10	10/100TX		Drop
11	10/100TX		Drop
12	10/100TX		Drop
.	.		.

ProCurve(config)# show filter 24

Traffic/Security Filters

Filter Type : Source Port

Source Port : 10

Dest	Port	Type	Action
-----	-----	+	-----
1	10/100TX		Drop
2	10/100TX		Drop
3	10/100TX		Drop
4	10/100TX		Drop
5	10/100TX		Drop
6	10/100TX		Drop
7	10/100TX		Forward
8	10/100TX		Drop
9	10/100TX		Drop
10	10/100TX		Drop
11	10/100TX		Drop
12	10/100TX		Drop
.	.		.

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

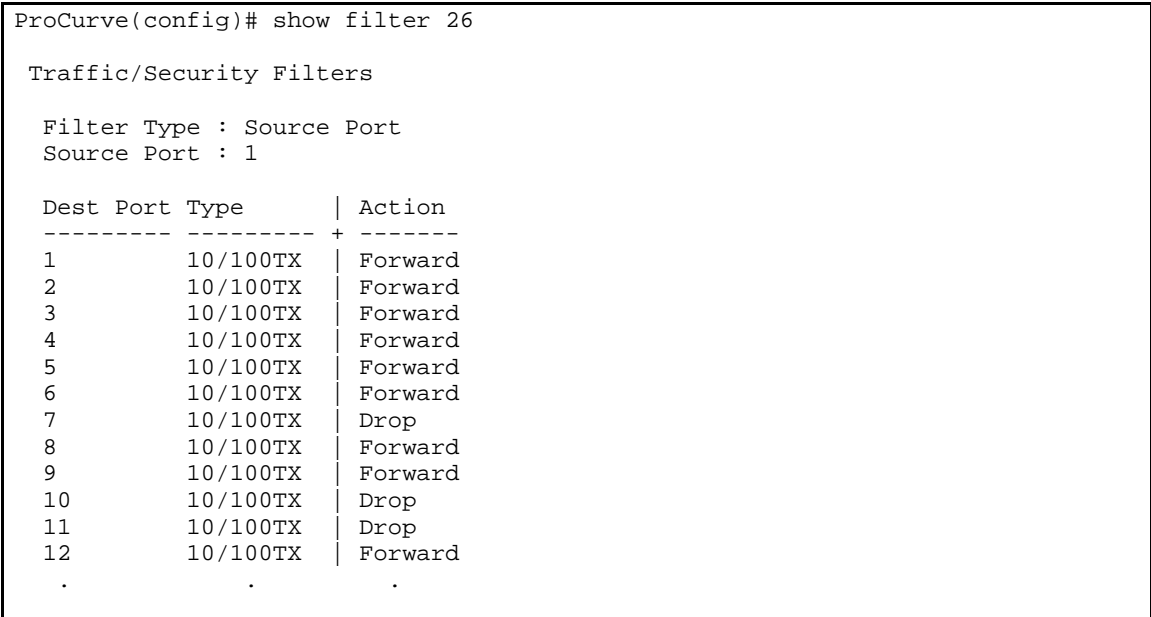


Figure 10-10. Show Filter Command Displaying Traffic from Internet

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.

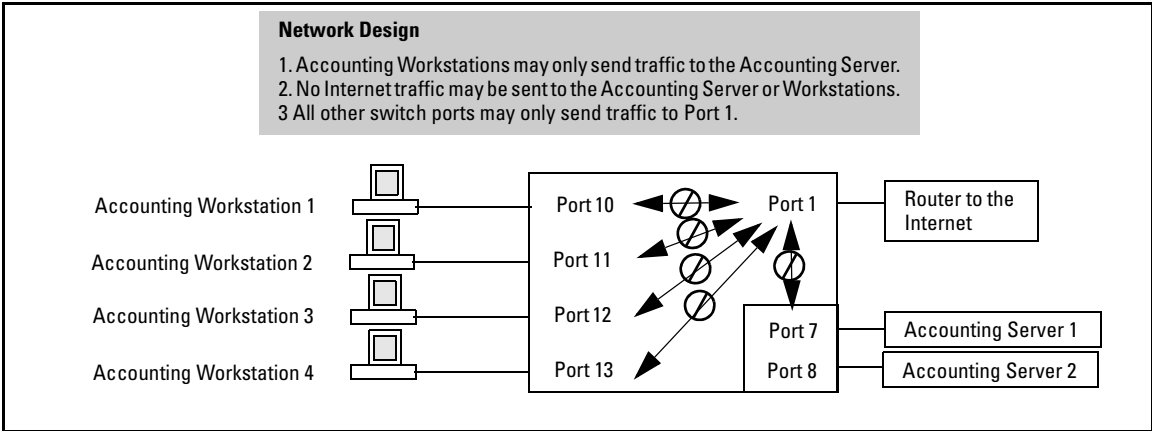


Figure 10-11. Expanded Network Configuration for Named Source-Port Filters Example

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.

```
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 10-12. Example of filter source-port Command Showing Actions

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 10-13. Removing Existing Source Port Filters

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

```
ProCurve(config)#
```

Figure 10-14. Displaying Traffic Filters

Configuring Port-Based and User-Based Access Control (802.1X)

Contents

Overview	11-3
Why Use Port-Based or User-Based Access Control?	11-3
General Features	11-3
User Authentication Methods	11-4
Terminology	11-6
General 802.1X Authenticator Operation	11-9
Example of the Authentication Process	11-9
VLAN Membership Priority	11-10
General Operating Rules and Notes	11-12
General Setup Procedure for 802.1X Access Control	11-14
Do These Steps Before You Configure 802.1X Operation	11-14
Overview: Configuring 802.1X Authentication on the Switch	11-16
Configuring Switch Ports as 802.1X Authenticators	11-17
1. Enable 802.1X Authentication on Selected Ports	11-18
2. Reconfigure Settings for Port-Access	11-20
3. Configure the 802.1X Authentication Method	11-24
4. Enter the RADIUS Host IP Address(es)	11-25
5. Enable 802.1X Authentication on the Switch	11-26
6. Optional: Reset Authenticator Operation	11-26
7. Optional: Configure 802.1X Controlled Directions	11-26
802.1X Open VLAN Mode	11-29
Introduction	11-29
VLAN Membership Priorities	11-30
Use Models for 802.1X Open VLAN Modes	11-31

Operating Rules for Authorized-Client and Unauthorized-Client VLANs	11-36
Setting Up and Configuring 802.1X Open VLAN Mode	11-40
802.1X Open VLAN Operating Notes	11-44
Option For Authenticator Ports: Configure Port-Security To Allow Only 802.1X-Authenticated Devices	11-45
Port-Security	11-46
Configuring Switch Ports To Operate As Suplicants for 802.1X Connections to Other Switches	11-47
Example	11-47
Supplicant Port Configuration	11-48
Displaying 802.1X Configuration, Statistics, and Counters	11-51
Show Commands for Port-Access Authenticator	11-51
Viewing 802.1X Open VLAN Mode Status	11-54
Show Commands for Port-Access Supplicant	11-57
How RADIUS/802.1X Authentication Affects VLAN Operation	11-58
VLAN Assignment on a Port	11-59
Operating Notes	11-59
Example of Untagged VLAN Assignment in a RADIUS-Based Authentication Session	11-61
Enabling the Use of GVRP-Learned Dynamic VLANs in Authentication Sessions	11-64
Operating Note	11-66
Messages Related to 802.1X Operation	11-67

Overview

Feature	Default	Menu	CLI	Web
Configuring Switch Ports as 802.1X Authenticators	Disabled	n/a	page 11-17	n/a
Configuring 802.1X Open VLAN Mode	Disabled	n/a	page 11-29	n/a
Configuring Switch Ports to Operate as 802.1X Supplicants	Disabled	n/a	page 11-47	n/a
Displaying 802.1X Configuration, Statistics, and Counters	n/a	n/a	page 11-51	n/a
How 802.1X Affects VLAN Operation	n/a	n/a	page 11-58	n/a
RADIUS Authentication and Accounting	Refer to chapter 5, "RADIUS Authentication and Accounting"			

Why Use Port-Based or User-Based Access Control?

Local Area Networks are often deployed in a way that allows unauthorized clients to attach to network devices, or allows unauthorized users to get access to unattended clients on a network. Also, the use of DHCP services and zero configuration make access to networking services easily available. This exposes the network to unauthorized use and malicious attacks. While access to the network should be made easy, uncontrolled and unauthorized access is usually not desirable. 802.1X simplifies security management by providing access control along with the ability to control user profiles from up to three RADIUS servers while allowing a given user to use the same entering valid user credentials for access from multiple points within the network.

General Features

802.1X on the switches covered in this guide 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 access 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).
 - User-Based access control option with support for up to 8 authenticated clients per-port.

- Port-Based access control option allowing authentication by a single client to open the port. This option does not force a client limit and, on a port opened by an authenticated client, allows unlimited client access without requiring further authentication.
- Supplicant implementation using CHAP authentication and independent user credentials on each port.
- The local operator password configured with the **password** command for management access to the switch is no longer accepted as an 802.1X authenticator credential. The **password port-access** command configures the local operator username and password used as 802.1X authentication credentials for access to the switch. The values configured can be stored in a configuration file using the **include-credentials** command. For information about the **password port-access** command, see “Do These Steps Before You Configure 802.1X Operation” on page 11-14.
- On-demand change of a port’s configured VLAN membership status to support the current client session.
- Session accounting with a RADIUS server, including the accounting update interval.
- Use of Show commands to display session counters.
- Support for concurrent use of 802.1X and either Web authentication or MAC authentication on the same port.
- For unauthenticated clients that do not have the necessary 802.1X supplicant software (or for other reasons related to unauthenticated clients), there is the option to configure an Unauthorized-Client VLAN. This mode allows you to assign unauthenticated clients to an isolated VLAN through which you can provide the necessary supplicant software and/or other services you want to extend to these clients.

User Authentication Methods

The switch offers two methods for using 802.1X access control. Generally, the “Port Based” method supports one 802.1X-authenticated client on a port, which opens the port to an unlimited number of clients. The “User-Based” method supports up to eight 802.1X-authenticated clients on a port. In both cases, there are operating details to be aware of that can influence your choice of methods.

802.1X User-Based Access Control

802.1X operation with access control on a per-user basis provides client-level security that allows LAN access to individual 802.1X clients (up to 8 per port), where each client gains access to the LAN by entering valid user credentials.

This operation improves security by opening a given port only to individually authenticated clients, while simultaneously blocking access to the same port for clients that cannot be authenticated. All sessions must use the same untagged VLAN. Also, an authenticated client can use any tagged VLAN memberships statically configured on the port, provided the client is configured to use the tagged VLAN memberships available on the port. (Note that the session total includes any sessions begun by the Web Authentication or MAC Authentication features covered in chapter 3.) For more information, refer to “Option For Authenticator Ports: Configure Port-Security To Allow Only 802.1X-Authenticated Devices” on page 11-45.

802.1X Port-Based Access Control

802.1X port-based access control provides port-level security that allows LAN access only on ports where a single 802.1X-capable client (supplicant) has entered authorized RADIUS user credentials. For reasons outlined below, this option is recommended for applications where only one client at a time can connect to the port. Using this option, the port processes all traffic as if it comes from the same client. Thus, in a topology where multiple clients can connect to the same port at the same time:

- If the first client authenticates and opens the port, and then another client authenticates, the port responds as if the original client has initiated a reauthentication. With multiple clients authenticating on the port, the RADIUS configuration response to the latest client authentication replaces any other configuration from an earlier client authentication. If all clients use the same configuration this should not be a problem. But if the RADIUS server responds with different configurations for different clients, then the last client authenticated will effectively lock out any previously authenticated client. When *any* client to authenticate closes its session, the port will also close and remain so until another client successfully authenticates.
- The most recent client authentication determines the untagged VLAN membership for the port. Also, any client able to use the port can access any tagged VLAN memberships statically configured on the port, provided the client is configured to use the available, tagged VLAN memberships.
- If the first client authenticates and opens the port, and then one or more other clients connect without trying to authenticate, then the port configuration as determined by the original RADIUS response remains unchanged and all such clients will have the same access as the authenticated client. When the authenticated client closes the session, the port will also be closed to any other, unauthenticated clients that may have also been using the port.

This operation unblocks the port while an authenticated client session is in progress. In topologies where simultaneous, multiple client access is possible this can allow unauthorized and unauthenticated access by another client while an authenticated client is using the port. If you want to allow only authenticated clients on the port, then user-based access control (page 11-4) should be used instead of port-based access control. Using the user-based method enables you to specify up to 8 authenticated clients.

Note

Port-Based 802.1X can operate concurrently with Web-Authentication or MAC-Authentication on the same port. However, this is not a commonly used application and is not generally recommended. For more information, refer to “Operating Note” on page 11-66.

Alternative To Using a RADIUS Server

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 user credential administration, and reduces security by limiting authentication to one Operator password set for all users.

Terminology

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

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

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

Authenticator: In ProCurve applications, a switch that requires a supplicant to provide the proper credentials before being allowed access to the network.

CHAP (MD5): Challenge Handshake Authentication Protocol.

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

User-Based Authentication: The 802.1X extension in the switches covered in this guide. In this operation, multiple clients on the same port must individually authenticate themselves.

Guest VLAN: See “Unauthorized-Client VLAN”.

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

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

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

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

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

Port-Based Authentication: In this operation, the first client on a port to authenticate itself unblocks the port for the duration of the client’s 802.1X-authenticated session. The switches covered in this guide use port-based authentication.

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

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

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

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

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

General 802.1X Authenticator Operation

This operation provides security on a point-to-point link between a 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 11-29.)

Example of the Authentication Process

Suppose that you have configured a port on the switch for 802.1X authentication operation, which blocks access to the LAN through that port. If you then connect an 802.1X-aware client (supplicant) to the port and attempt to log on:

1. The switch responds with an identity request.
2. The client responds with a user name that uniquely defines this request for the client.
3. The switch responds in one of the following ways:
 - If 802.1X 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 on the switch is configured for local authentication, then:
 - i. The switch compares the client's credentials to the username and password configured in the switch (Operator 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 for that client.

Note

The switches covered in this guide can use either 802.1X port-based authentication or 802.1X user-based authentication. For more information, refer to “User Authentication Methods” on page 11-4.

VLAN Membership Priority

Following client authentication, an 802.1X port resumes membership in any tagged VLANs for which it is already assigned in the switch configuration. The port also becomes an untagged member of one VLAN according to the following order of options:

- a. **1st Priority:** The port joins a VLAN to which it has been assigned by a RADIUS server during client authentication.
- b. **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.
- c. **3rd Priority:** If the port does not have an Authorized-Client VLAN configured, but does have a static, untagged VLAN membership in its configuration, then the switch assigns the port to this VLAN.

A port assigned to a VLAN by an Authorized-Client VLAN configuration (or a RADIUS server) will be an untagged member of the VLAN for the duration of the authenticated session. This applies even if the port is also configured in the switch as a tagged member of the same VLAN.

Note

On the switches covered in this guide, using the same port for both RADIUS-assigned clients and clients using a configured, Authorized-Client VLAN is not recommended. This is because doing so can result in authenticated clients with mutually exclusive VLAN priorities, which means that some authenticated clients can be denied access to the port. Refer to figure 11-1 on page 11-11.

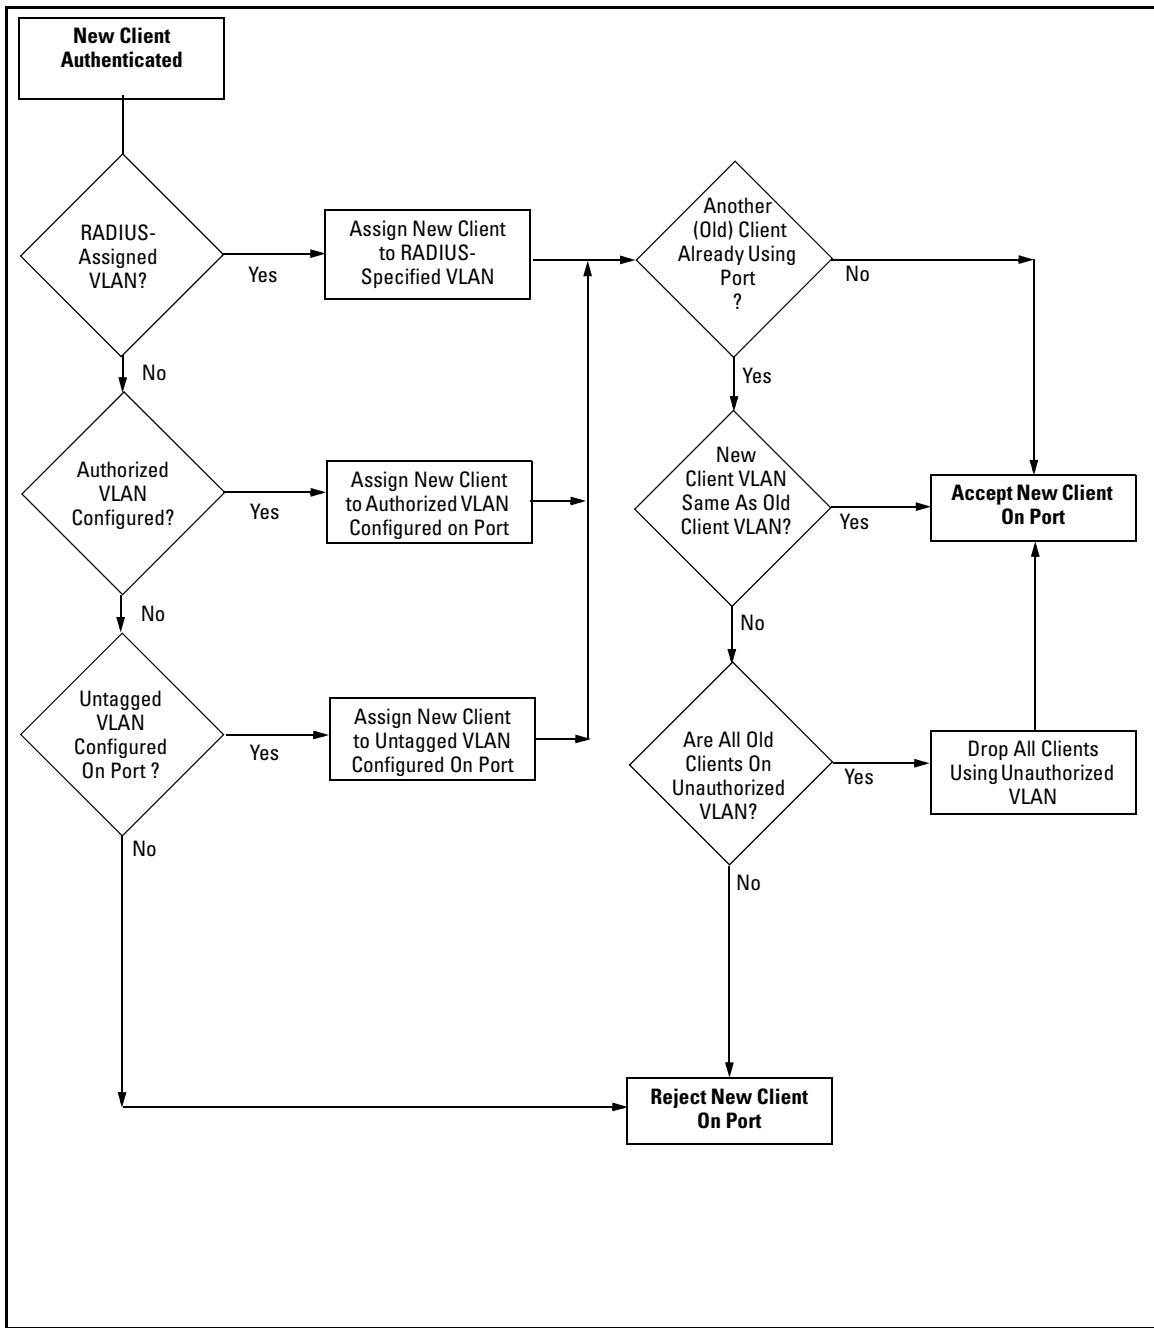


Figure 11-1. Priority of VLAN Assignment for an Authenticated Client

General Operating Rules and Notes

- In the user-based mode, when there is an authenticated client on a port, the following traffic movement is allowed:
 - Multicast and broadcast traffic is allowed on the port.
 - Unicast traffic to authenticated clients on the port is allowed.
 - All traffic from authenticated clients on the port is allowed.
- When a port on the switch is configured as either an authenticator or supplicant and is connected to another device, rebooting the switch causes a re-authentication of the link.
- Using user-based 802.1X authentication, when a port on the switch is configured as an authenticator the port allows only authenticated clients up to the currently configured client limit.

For clients that do not have the proper 802.1X supplicant software, the optional 802.1X Open VLAN mode can be used to open a path for downloading 802.1X supplicant software to a client or to provide other services for unauthenticated clients. Refer to “802.1X Open VLAN Mode” on page 11-29.)

- Using port-based 802.1X authentication, When a port on the switch is configured as an authenticator, one authenticated client opens the port. Other clients that are not running an 802.1X supplicant application can have access to the switch and network through the opened port. If another client uses an 802.1X supplicant application to access the opened port, then a re-authentication occurs using the RADIUS configuration response for the latest client to authenticate. To control access by all clients, use the user-based method.
- Where a switch port is configured with user-based authentication to accept multiple 802.1X (and/or Web- or MAC-Authentication) client sessions, all authenticated clients must use the same port-based, untagged VLAN membership assigned for the earliest, currently active client session. Thus, on a port where one or more authenticated client sessions are already running, all such clients will be on the same untagged VLAN. If a RADIUS server subsequently authenticates a new client, but attempts to re-assign the port to a different, untagged VLAN than the one already in use for the previously existing, authenticated client sessions, the connection for the new client will fail. For more on this topic, refer to “802.1X Open VLAN Mode” on page 11-29. (Note that if the port is statically configured with any tagged VLAN memberships, any authenticated client configured to use these tagged VLANs will have access to them.)

- 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.
- 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 allow authentication of the supplicant. 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 will allow the supplicant to re-authenticate.
- 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.
- Meshing is not supported on ports configured for 802.1X port-access security.
- A port can be configured as an authenticator *or* an 802.1X supplicant, or both. Some configuration instances block traffic flow or allow traffic to flow without authentication. Refer to “Configuring Switch Ports To Operate As Supplicants for 802.1X Connections to Other Switches” on page 11-47.
- To help maintain security, 802.1X and LACP cannot both be enabled on the same port. If you try to configure 802.1X on a port already configured for LACP (or the reverse) you will see a message similar to the following:

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

General Setup Procedure for 802.1X Access Control

Do These Steps Before You Configure 802.1X Operation

1. Configure a local username and password on the switch for both the Operator (login) and Manager (enable) access levels. (While this may or may not be required for your 802.1X configuration, ProCurve recommends that you use a local username and password pair at least until your other security measures are in place.)

For switches covered in this guide, the local operator password configured with the password command is not accepted as an 802.1X authenticator credential. The port-access command is used to configure the operator username and password that are used as 802.1X credentials for network access to the switch. 802.1X network access is not allowed unless a password has been configured using the **password port-access** command.

Syntax: password port-access [user-name <name>] <password>

Configured the operator username and password used to access the network through 802.1X authentication.

user-name <name>

*Operator username (text string) used only for local authentication of 802.1X clients. This value is different from the local operator username configured with the **password** command for management access.*

<password>

*Operator password (text string) used only for local authentication of 802.1X clients. This value is different from the local operator password configured with the **password** command for management access.*

Figure 11-2 shows how to configure a local operator password for 802.1X access.

```
ProCurve(config)# password port-access user-name Jim secret3
```

Figure 11-2. Example of the Password Port-Access Command

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 page 11-18.)
3. Determine whether to use user-based access control (page 11-4) or port-based access control (page 11-5).
4. Determine whether to use the optional 802.1X Open VLAN mode for clients that are not 802.1X-aware; that is, for clients that are not running 802.1X supplicant software. (This will require you to provide downloadable software that the client can use to enable an authentication session.) For more on this topic, refer to “802.1X Open VLAN Mode” on page 11-29.
5. For any port you want to operate as a supplicant, determine the user credentials. You can either use the same credentials for each port or use unique credentials for individual ports or subgroups of ports. (This can also be the same local username/password pair that you assign to the switch.)
6. Unless you are using only the switch’s local username and password for 802.1X authentication, configure at least one RADIUS server to authenticate access requests coming through the ports on the switch from external supplicants (including switch ports operating as 802.1X supplicants). You can use up to three RADIUS servers for authentication; one primary and two backups. Refer to the documentation provided with your RADIUS application.

Overview: Configuring 802.1X Authentication on the Switch

This section outlines the steps for configuring 802.1X on the switch. For detailed information on each step, refer to the following:

- “802.1X User-Based Access Control” on page 11-4
 - “802.1X Port-Based Access Control” on page 11-5
 - “Configuring Switch Ports To Operate As Supplicants for 802.1X Connections to Other Switches” on page 11-47.
1. Enable 802.1X user-based or port-based 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 11-18.
 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 11-29.
 3. Configure the 802.1X authentication type. Options include:
 - Local Operator username and password (using the **password port-access** command).
 - 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.Refer to page 11-24.
 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 11-25.
 5. Enable 802.1X authentication on the switch. Refer to “1. Enable 802.1X Authentication on Selected Ports” on page 11-18.
 6. Test both the authorized and unauthorized access to your system to ensure that the 802.1X authentication works properly on the ports you have configured for port-access.

Note

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

7. If you are using Port Security on the switch, configure the switch to allow only 802.1X access on ports configured for 802.1X operation, and (if desired) the action to take if an unauthorized device attempts access through an 802.1X port. Refer to page 11-45.
8. If you want a port on the switch to operate as a supplicant on 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 11-47.)

Configuring Switch Ports as 802.1X Authenticators

802.1X Authentication Commands	Page
[no] aaa port-access authenticator < <i>port-list</i> >	11-18
[auth-vid clear-statistics client-limit control max-requests initialize logoff-period quiet-period server-timeout reauthenticate reauth-period supplicant-timeout tx-period unauth-period unauth-vid]	11-18
aaa authentication port-access < local eap-radius chap-radius >	11-24
[no] aaa port-access authenticator active	11-17
aaa port-access < <i>port-list</i> > controlled-direction < both in >	11-26
[no] port-security [ethernet] < <i>port-list</i> > learn-mode port-access	11-45
802.1X Open VLAN Mode Commands	11-29
802.1X Supplicant Commands	11-47
802.1X-Related Show Commands	11-51
RADIUS server configuration	11-25

1. Enable 802.1X Authentication on Selected Ports

This task configures the individual ports you want to operate as 802.1X authenticators for point-to-point links to 802.1X-aware clients or switches, and consists of two steps:

- A. Enable the selected ports as authenticators.
- B. Specify either user-based or port-based 802.1X authentication.

(Actual 802.1X operation does not commence until you perform step 5 on page 11-26 to activate 802.1X authentication on the switch.)

Note

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

A. Enable the Selected Ports as Authenticators and Enable the (Default) Port-Based Authentication

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

*Enables specified ports to operate as 802.1X authenticators and enables port-based authentication. (To enable user-based authentication, execute this command first, and then execute the client-limit < port-list > version of this command described in the next section.) The **no** form of the command removes 802.1X authentication from < port-list >. To activate configured 802.1X operation, you must enable 802.1X authentication. Refer to “5. Enable 802.1X Authentication on the switch” on page 11-26.*

B. Specify User-Based Authentication or Return to Port-Based Authentication

User-Based 802.1X Authentication.

Syntax: `aaa port-access authenticator client-limit < port-list > < 1 - 8>`

*Used after executing **aaa port-access authenticator < port-list >** (above) to convert authentication from port-based to user-based. Specifies user-based 802.1X authentication and the maximum number of 802.1X-authenticated client sessions allowed on each of the ports in **< port-list >**. If a port currently has no authenticated client sessions, the next authenticated client session the port accepts determines the untagged VLAN membership to which the port is assigned during the session. If another client session begins later on the same port while an earlier session is active, the later session will be on the same untagged VLAN membership as the earlier session.*

Note: *Because a switch allows 802.1X authentication and Web or MAC authentication to co-exist on the same port, the sum of authenticated client sessions allowed on a given port for both 802.1X and either Web- or MAC-authentication cannot exceed 32.*

Port-Based 802.1X Authentication.

`no aaa port-access authenticator client-limit`

*Used to convert a port from user-based authentication to port-based authentication, which is the default setting for ports on which authentication is enabled. (Executing **aaa port-access authenticator < port-list >** enables 802.1X authentication on **< port-list >** and enables port-based authentication—page 11-18.) If a port currently has no authenticated client sessions, the next authenticated client session the port accepts determines the untagged VLAN membership to which the port is assigned during the session. If another authenticated client session begins later on the same port while an earlier session is active, the later session replaces the currently active session and will be on the untagged VLAN membership specified by the RADIUS server for the later session.*

Example: Configuring User-Based 802.1X Authentication

This example enables ports A10-A12 to operate as authenticators, and then configures the ports for user-based authentication.

```
ProCurve(config)# aaa port-access authenticator a10-A12
ProCurve(config)# aaa port-access authenticator a10-A12 client-limit 4
```

Figure 11-3. Example of Configuring User-Based 802.1X Authentication

Example: Configuring Port-Based 802.1X Authentication

This example enables ports A13-A15 to operate as authenticators, and then configures the ports for port-based authentication.

```
ProCurve(config)# aaa port-access authenticator al3-al5
ProCurve(config)# no aaa port-access authenticator al3-al5 client-limit
```

Figure 11-4. Example of Configuring Port-Based 802.1X Authentication

2. Reconfigure Settings for Port-Access

The commands in this section are initially set by default and can be reconfigured as needed.

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

Controls authentication mode on the specified port:

authorized: Also termed “Force Authorized”. Gives access to a device connected to the port. In this case, the device does not have to provide 802.1X credentials or support 802.1X authentication. (You can still configure console, Telnet, or SSH security on the port.)

auto (the default): The device connected to the port must support 802.1X authentication and provide valid credentials to get network access. (Optional: You can use the Open VLAN mode to provide a path for clients without 802.1X supplicant software to down-load this software and begin the authentication process. Refer to “802.1X Open VLAN Mode” on page 11-29.)

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

[quiet-period < 0 - 65535 >]

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

[tx-period < 0 - 65535 >]

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

[supplicant-timeout < 1 - 300 >]

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

[server-timeout < 1 - 300 >]

*Sets the period of time the switch waits for a server response to an authentication request. If 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)*

—Continued—

[quiet-period < 0 - 65535 >]

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

[tx-period < 0 - 65535 >]

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

[supplicant-timeout < 1 - 300 >]

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

[server-timeout < 1 - 300 >]

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

[max-requests < 1 - 10 >]

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

—Continued—

[reauth-period < 0 - 9999999 >]

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

[unauth-vid < vlan-id >]

Configures an existing static VLAN to be the Unauthorized-Client VLAN. This enables you to provide a path for clients without supplicant software to download the software and begin an authentication session. Refer to “802.1X Open VLAN Mode” on page 11-29.

[logoff-period]< 1 - 999999999 >

Configures the period of time the switch waits for client activity before removing an inactive client from the port. (Default: 300 seconds)

[unauth-period < 0-255 >]

Specifies a delay in seconds for placing a port on the Unauthorized-Client VLAN. This delay allows more time for a client with 802.1X supplicant capability to initiate an authentication session. If a connected client does not initiate a session before the timer expires, the port is assigned to the Unauthenticated-Client VLAN. (Default: 0 seconds)

[auth-vid < vid >]

Configures an existing, static VLAN to be the Authorized-Client VLAN. Refer to “802.1X Open VLAN Mode” on page 11-29.

3. Configure the 802.1X Authentication Method

This task specifies how the switch authenticates 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.

*For switches covered in this guide, you must use the **password port-access** command to configure the operator username and password for 802.1X access. See "General Setup Procedure for 802.1X Access Control" on page 11-14 for more information.*

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

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

<none | local>

Provides options for secondary authentication.

*The **authorized** option allows users unconditional access to the network when the primary authentication method fails. See **Caution** below.*

*Default: **None***

Caution

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

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

Configuration command for EAP-RADIUS 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	Local	None	Local	None
Port-Access	EapRadius			
Webui	Local	None	Local	None
SSH	Local	None	Local	None
Web-Auth	ChapRadius			
MAC-Auth	ChapRadius			

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

Figure 11-5. Example of 802.1X (Port-Access) Authentication

4. Enter the RADIUS Host IP Address(es)

If you select either **eap-radius** or **chap-radius** for the authentication method, configure the switch to use 1, 2, or 3 RADIUS servers for authentication. The following syntax shows the basic commands. For coverage of all commands related to RADIUS server configuration, refer to chapter 5, “RADIUS Authentication and Accounting”.

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.

Syntax: radius-server key < global key-string >

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

5. Enable 802.1X Authentication on the Switch

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

Syntax: aaa port-access authenticator active

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

6. Optional: Reset Authenticator Operation

While 802.1X authentication is operating, you can use the following **aaa port-access authenticator** commands to reset 802.1X authentication and statistics on specified ports.

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

[initialize]

*On the specified ports, blocks inbound and outbound traffic and restarts the 802.1X authentication process. This happens only on ports configured with **control auto** and actively operating as 802.1X authenticators.*

[reauthenticate]

On the specified ports, forces reauthentication (unless the authenticator is in “HELD” state).

[clear-statistics]

On the specified ports, clears authenticator statistics counters.

7. Optional: Configure 802.1X Controlled Directions

After you enable 802.1X authentication on specified ports, you can use the **aaa port-access controlled-directions** command to configure how a port transmits traffic before it successfully authenticates a client and enters the authenticated state.

As documented in the IEEE 802.1X standard, an 802.1X-aware port that is unauthenticated can control traffic in either of the following ways:

- In both ingress and egress directions by disabling both the reception of incoming frames and transmission of outgoing frames
- Only in the ingress direction by disabling only the reception of incoming frames.

Prerequisite. As documented in the IEEE 802.1X standard, the disabling of incoming traffic and transmission of outgoing traffic on an 802.1X-aware egress port in an unauthenticated state (using the **aaa port-access controlled-directions in** command) is supported only if:

- The port is configured as an edge port in the network using the **spanning-tree edge-port** command.
- The 802.1s Multiple Spanning Tree Protocol (MSTP) or 802.1w Rapid Spanning Tree Protocol (RSTP) is enabled on the switch. MSTP and RSTP improve resource utilization while maintaining a loop-free network.

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

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

both (default): *Incoming and outgoing traffic is blocked on an 802.1X-aware port before authentication occurs.*

in: *Incoming traffic is blocked on an 802.1X-aware port before authentication occurs. Outgoing traffic with unknown destination addresses is flooded on unauthenticated 802.1X-aware ports.*

Wake-on-LAN Traffic

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

The **aaa port-access controlled-direction in** command allows Wake-on-LAN traffic to be transmitted on an 802.1X-aware egress port that has not yet transitioned to the 802.1X authenticated state; the **controlled-direction both** setting prevents Wake-on-LAN traffic to be transmitted on an 802.1X-aware egress port until authentication occurs.

Note

Although the **controlled-direction in** setting allows Wake-on-LAN traffic to traverse the switch through unauthenticated 802.1X-aware egress ports, it does not guarantee that the Wake-on-LAN packets will arrive at their destination. For example, firewall rules on other network devices and VLAN rules may prevent these packets from traversing the network.

Operating Notes

- Using the **aaa port-access controlled-directions in** command, you can enable the transmission of Wake-on-LAN traffic on unauthenticated egress ports that are configured for any of the following port-based security features:
 - 802.1X authentication
 - MAC authentication
 - Web authentication

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

For information about how to configure and use MAC and Web authentication, refer to chapter 3, “Web and MAC Authentication”.

- To display the currently configured 802.1X Controlled Directions value, enter the **show port-access authenticator config** command as shown in Figure 11-9.
- When an 802.1X-authenticated port is configured with the **controlled-directions in** setting, eavesdrop prevention is not supported on the port.

Example: Configuring 802.1X Controlled Directions

The following example shows how to enable the transmission of Wake-on-LAN traffic in the egress direction on an 802.1X-aware port before it transitions to the 802.1X authenticated state and successfully authenticates a client device.

```
ProCurve(config)# aaa port-access authenticator a10
ProCurve(config)# aaa authentication port-access eap-radius
ProCurve(config)# aaa port-access authenticator active
ProCurve(config)# aaa port-access a10 controlled-directions in
```

Figure 11-6. Example of Configuring 802.1X Controlled Directions

802.1X Open VLAN Mode

802.1X Authentication Commands	page 11-17
802.1X Supplicant Commands	page 11-48
802.1X Open VLAN Mode Commands	
[no] aaa port-access authenticator < <i>port-list</i> >	page 11-43
[auth-vid < <i>vlan-id</i> >]	
[unauth-vid < <i>vlan-id</i> >]	
802.1X-Related Show Commands	page 11-51
RADIUS server configuration	pages 11-25

Introduction

This section describes how to use the 802.1X Open VLAN mode to provide a path for clients that need to acquire 802.1X supplicant software before proceeding with the authentication process. The Open VLAN mode involves options for configuring unauthorized-client and authorized-client VLANs on ports configured as 802.1X authenticators.

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

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

The 802.1X Open VLAN mode solves this problem by temporarily suspending the port’s static VLAN memberships and placing the port in a designated *Unauthorized-Client VLAN* (sometimes termed a *guest VLAN*). In this state the client can proceed with initialization services, such as acquiring IP addressing and 802.1X client software, and starting the authentication process.

Note

On ports configured to allow multiple sessions using 802.1X user-based access control, all clients must use the same untagged VLAN. On a given port where there are no currently active, authenticated clients, the first *authenticated* client determines the untagged VLAN in which the port will operate for all subsequent, overlapping client sessions.

If the switch operates in an environment where some valid clients will not be running 802.1X supplicant software and need to download it from your network. Then, because such clients would need to use the Unauthorized-Client VLAN and authenticated clients would be using a different VLAN (for security reasons), allowing multiple clients on an 802.1X port can result in blocking some or all clients needing to use the Unauthorized-Client VLAN.

On ports configured for port-based 802.1X access control, if multiple clients try to authenticate on the same port, the most recently authenticated client determines the untagged VLAN membership for that port. Clients that connect without trying to authenticate will have access to the untagged VLAN membership that is currently assigned to the port.

VLAN Membership Priorities

Following client authentication, an 802.1X port resumes membership in any tagged VLANs for which it is already assigned in the switch configuration. The port also becomes an untagged member of one VLAN according to the following order of options:

- a. **1st Priority:** The port joins a VLAN to which it has been assigned by a RADIUS server during client authentication.
- b. **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.
- c. **3rd Priority:** If the port does not have an Authorized-Client VLAN configured, but does have a static, untagged VLAN membership in its configuration, then the switch assigns the port to this VLAN.

A port assigned to a VLAN by an Authorized-Client VLAN configuration (or a RADIUS server) will be an untagged member of the VLAN for the duration of the authenticated session. This applies even if the port is also configured in the switch as a tagged member of the same VLAN.

Note

After client authentication, the port resumes membership in any tagged VLANs for which it is configured. If the port is a tagged member of a VLAN used for 1 or 2 listed above, then it also 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 or instead of 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 port-based 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 it is configured. See the "Note", above.

Table 11-1. 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 without 802.1X supplicant capability, 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. <p>Note for a Port Configured To Allow Multiple Client Sessions: If any previously authenticated clients are using a port assigned to a VLAN other than the Unauthorized-Client VLAN, then a later client that is not running 802.1X supplicant software is blocked on the port until all other, authenticated clients on the port have disconnected.</p>

802.1X Per-Port Configuration	Port Response
Authorized-Client VLAN	<ul style="list-style-type: none">After client authentication, the port drops membership in the Unauthorized-Client VLAN and becomes an untagged member of this VLAN. <p>Notes: If the client is running an 802.1X supplicant application when the authentication session begins, and is able to authenticate itself before the switch assigns the port to the Unauthorized-Client VLAN, then the port does not become a member of the Unauthorized-Client VLAN. On the switches covered in this guide, you can use the unauth-period command—page 11-23—to delay moving the port into the Unauthorized-Client VLAN.</p> <p>If RADIUS authentication assigns a VLAN and there are no other authenticated clients on the port, then the port becomes a member of the RADIUS-assigned VLAN—instead of the Authorized-Client VLAN—while the client is connected.</p> <ul style="list-style-type: none">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.If the port is statically configured as a tagged member of a VLAN, 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.

802.1X Per-Port Configuration	Port Response
Open VLAN Mode with Only an Unauthorized-Client VLAN Configured:	
	<ul style="list-style-type: none">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. <p>Note: If RADIUS authentication assigns the port to a VLAN, this assignment overrides any statically configured, untagged VLAN membership on the port (while the client is connected).</p> <ul style="list-style-type: none">If the port is statically configured as a tagged member of a 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. 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 <i>untagged</i> member of that VLAN for the duration of the client connection. <p>Note for a Port Configured To Allow Multiple Client Sessions: If any previously authenticated clients are using a port assigned to a VLAN other than the Unauthorized-Client VLAN (such as a RADIUS-assigned VLAN), then a later client that is not running 802.1X supplicant software is blocked on the port until all other, authenticated clients on the port have disconnected. Refer to figure 11-1 on page 11-11.</p>

802.1X Per-Port Configuration	Port Response
Open VLAN Mode with Only an Authorized-Client VLAN Configured:	
	<ul style="list-style-type: none">• Port automatically blocks a client that cannot initiate an authentication session.• If the client successfully completes an authentication session, the port becomes an untagged member of this VLAN.• If the port is statically configured as a tagged member of any other 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 <i>untagged</i> member of that VLAN for the duration of the client connection.
	Note: An authorized-client VLAN configuration can be overridden by a RADIUS authentication that assigns a VLAN. (Refer to figure 11-1 on page 11-11.)

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 membership in both VLANs is untagged, and the switch allows only one untagged, port-based 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 11-31.)
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. In the case of the multiple clients allowed on switches covered in this guide, the first client to authenticate determines the untagged VLAN membership for the port until all clients have disconnected. Any other clients that cannot operate in that VLAN are blocked at that point.• Port membership in a VLAN assigned to operate as the Authorized-Client VLAN ends when the client disconnects from the port. If a VLAN assignment from a RADIUS server is used instead, the same rule applies. In the case of the multiple clients allowed on switches, the port maintains the same VLAN as long as there is any authenticated client using the VLAN. When the last client disconnects, then the port reverts to only the VLAN(s) for which it is statically configured as a member.

Condition	Rule
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 any other VLANs.) If the client disconnects, the port leaves the Unauthorized-Client VLAN and re-acquires membership in all the statically configured VLANs to which it belongs. If the client becomes authenticated, the port leaves the Unauthorized-Client VLAN and joins the appropriate VLAN. (Refer to “VLAN Membership Priorities” on page 11-30. In the case of the multiple clients allowed on switches, if an authenticated client is already using the port for a different VLAN, then any other unauthenticated clients needing to use the Unauthorized-Client VLAN are blocked.
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 11-31.) <p>Note: This rule assumes:</p> <ul style="list-style-type: none"> No alternate VLAN has been assigned by a RADIUS server. No other authenticated clients are already using the port.
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-client VLAN 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 This rule assumes no other authenticated clients are already using the port on a different VLAN.	<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.)</p>

Configuring Port-Based and User-Based Access Control (802.1X)

802.1X Open VLAN Mode

Condition	Rule
Effect of RADIUS-assigned VLAN This rule assumes no other authenticated clients are already using the port on a different VLAN.	The port joins the RADIUS-assigned VLAN as an untagged member.
IP Addressing for a Client Connected to a Port Configured for 802.x Open VLAN Mode	A client can either acquire an IP address from a DHCP server or use a manually configured IP address before connecting to the switch.
802.1X Supplicant Software for a Client Connected to a Port Configured for 802.1X Open VLAN Mode	A friendly client, without 802.1X supplicant software, connecting to an authenticator port must be able to download this software from the Unauthorized-Client VLAN before authentication can begin.
Switch with a Port Configured To Allow Multiple Authorized-Client Sessions	<p>When a new client is authenticated on a given port:</p> <ul style="list-style-type: none">• If no other clients are authenticated on that port, then the port joins one VLAN in the following order of precedence:<ol style="list-style-type: none">a. A RADIUS-assigned VLAN, if configured.b. An Authenticated-Client VLAN, if configured.c. A static, port-based VLAN to which the port belongs as an untagged member.d. Any VLAN(s) to which the port is configured as a tagged member (provided that the client can operate in that VLAN).• If another client is already authenticated on the port, then the port is already assigned to a VLAN for the previously-existing client session, and the new client must operate in this same VLAN, regardless of other factors. (This means that a client without 802.1X client authentication software cannot access a configured, Unauthenticated-Client VLAN if another, authenticated client is already using the port.)

Condition	Rule
Note: Limitation on Using an Unauthorized-Client VLAN on an 802.1X Port Configured to Allow Multiple-Client Access	You can optionally enable switches to allow up to 8 clients per-port. The Unauthorized-Client VLAN feature can operate on an 802.1X-configured port regardless of how many clients the port is configured to support. However, all clients on the same port must operate through the same untagged VLAN membership. This means that any client accessing a given port must be able to authenticate and operate on the same VLAN as any other previously authenticated clients that are currently using the port. Thus, an Unauthorized-Client VLAN configured on a switch port that allows multiple 802.1X clients cannot be used if there is already an authenticated client using the port on another VLAN. Also, a client using the Unauthenticated-Client VLAN will be blocked when another client becomes authenticated on the port. For this reason, the best utilization of the Unauthorized-Client VLAN feature is in instances where only one client is allowed per-port. Otherwise, unauthenticated clients are subject to being blocked at any time by authenticated clients using a different VLAN. (Using the same VLAN for authenticated and unauthenticated clients can create a security risk and is not recommended.)

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

Setting Up and Configuring 802.1X Open VLAN Mode

Preparation. This section assumes use of both the Unauthorized-Client and Authorized-Client VLANs. Refer to Table 11-1 on page 11-32 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 port-control parameter is set to **auto** (the default). (Refer to “1. Enable 802.1X Authentication on Selected Ports” on page 11-18.) 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 < 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.

Syntax: 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-Authenticated Devices” on page 11-45.

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 11-40.

Syntax: aaa port-access authenticator < 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 11-54.

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”, 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. (If there is no Authorized-Client or RADIUS-assigned VLAN, then an authenticated client without tagged VLAN capability can access only a statically configured, untagged VLAN on that port.)
- When a client’s authentication attempt on an Unauthorized-Client VLAN fails, the port remains a member of the Unauthorized-Client VLAN until the client disconnects from the port.
- During an authentication session on a port in 802.1X Open VLAN mode, if RADIUS specifies membership in an untagged VLAN, this assignment overrides port membership in the Authorized-Client VLAN. If there is no Authorized-Client VLAN configured, then the RADIUS assignment overrides any untagged VLAN for which the port is statically configured.
- If the only authenticated client on a port loses authentication during a session in 802.1X Open VLAN mode, the port VLAN membership reverts back to the Unauthorized-Client VLAN. If there is no Unauthorized-Client VLAN configured, then the client loses access to the port until it can

reauthenticate itself. If there are multiple clients authenticated on the port, if one client loses access and attempts to re-authenticate, that client will be handled as a new client on the port.

- The first client to authenticate on a port configured to support multiple clients will determine the port’s VLAN membership for any subsequent clients that authenticate while an active session is already in effect.

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

If 802.1X authentication is disabled on a port or set to **authorized** (Force Authorize), the port can allow access to a non-authenticated client. Port-Security operates with 802.1X authentication only if the selected ports are configured as 802.1X *with* the **control** mode in the port-access authenticator command set to **auto** (the default setting). For example, if port A10 was at a non-default 802.1X setting and you wanted to configure it to support the port-security option, you would use the following **aaa port-access** command:

ProCurve(config)# aaa port-access authenticator a10 control auto
ProCurve(config)# show port-access authenticator a10 config

Control mode
required for Port-
Security Support

Port Access Authenticator Configuration

Port-access authenticator activated [No] : No

Port	Re-auth Period	Access Control	Max Requests	Quiet Period	TX Timeout	Supplicant Timeout	Server Timeout
A10	No	Auto	2	60	30	30	30

Figure 11-7. Port-Access Support for Port-Security Operation

Port-Security

Note

If 802.1X port-access is configured on a given port, then port-security **learn-mode** for that port must be set to either **continuous** (the default) or **port-access**.

In addition to the above, to use port-security on an authenticator port (chapter 12), use the per-port **client-limit** option to control how many MAC addresses of 802.1X-authenticated devices the port is allowed to learn. (Using **client-limit** sets 802.1X to user-based operation on the specified ports.) When this limit is reached, no further devices can be authenticated until a currently authenticated device disconnects and the current delay period or logoff period has expired.

Configure the port access type.

Syntax: aaa port-access auth < port-list > client-limit < 1 - 8 >

Configures user-based 802.1X authentication on the specified ports and sets the number of authenticated devices the port is allowed to learn. For more on this command, refer to “Configuring Switch Ports as 802.1X Authenticators” on page 11-17.)

— Or —

no aaa port-access auth < port-list > client-limit

Configures port-based 802.1X authentication on the specified ports, which opens the port. (Refer to “User Authentication Methods” on page 11-4.)

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

802.1X Authentication Commands	page 11-17
802.1X Supplicant Commands	
[no] aaa port-access < supplicant < [ethernet] < port-list >	page 11-48
[auth-timeout held-period start-period max-start initialize identity secret clear-statistics]	page 11-48
802.1X-Related Show Commands	page 11-51
RADIUS server configuration	pages 11-25

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

Example

Suppose that you want to connect two switches, where:

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




Figure 11-8. 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 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, the authenticator compares the switch “A” response to its local username and password.
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.

Supplicant Port Configuration

Enabling a Switch Port as a Supplicant. You can configure a switch port as a supplicant for a point-to-point link to an 802.1X-aware port on another switch. *Configure the port as a supplicant before configuring any supplicant-related parameters.*

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

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

Configuring a Supplicant Switch Port. You must enable supplicant operation on a port before changing 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 credentials 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 due 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. (Default: Null.)*

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

`[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.

`[auth-timeout < 1 - 300 >]`

*Sets the delay period the port waits to receive a challenge from the authenticator. If the request times out, the port sends another 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 11-47 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 delay 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 11-17
802.1X Supplicant Commands	page 11-47
802.1X Open VLAN Mode Commands	page 11-29
802.1X-Related Show Commands	
show port-access authenticator	below
show port-access supplicant	page 11-57
Details of 802.1X Mode Status Listings	page 11-54
RADIUS server configuration	pages 11-25

Show Commands for Port-Access Authenticator

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

- *Without [config | statistics | session-counters | vlan] [<port-list>], displays whether port-access authenticator is active (**Yes** or **No**) and the status of all ports configured for 802.1X authentication.*
- *With <port-list> only, same as above, but only for the specified ports. Does not display data for a specified port that is not enabled as an authenticator.*
- *With [config | statistics | session-counters | vlan] [<port-list>], 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 | vlan] only, displays the [config | statistics | session-counters] data for all ports enabled as authenticators.*

For more information on the [config | statistics | session-counters | vlan] options, refer to the next section of this table.

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

config [< port-list>]

Shows:

- Whether port-access authenticator is active
- The 802.1X configuration settings of ports configured as 802.1X authenticators (For a description of each setting, refer to the syntax descriptions in “2. Reconfigure Settings for Port-Access” on page 11-20. Use **show running** to view the current **client-limit** configuration available for switches.)

Without <port-list>, the command lists ports configured as 802.1X port-access authenticators. Does not display data for a port not enabled as an authenticator.

statistics [< 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 [< port-list>]

Shows whether port-access authenticator is active, and includes 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 used in its response packet. (For the switch, this is the **identity** setting included in the **supplicant** command—page 11-48.) Does not display data for a specified port that is not an authenticator.*

vlan [< port-list>]

Shows per-port:

- The Access Control setting (**control** command on page 11-20)
- Unauth-VLAN ID (if any)
- Auth-VLAN ID (if any)


```
ProCurve(config)# show port-access authenticator config
```

Port Access Authenticator Configuration

Port-access authenticator activated [No] : No

Port	Re-auth Period	Access Control	Max Reqs	Quiet Period	TX Timeout	Supplicant Timeout	Server Timeout	Cntrl Dir
1	No	Auto	2	60	30	30	30	both
2	No	Auto	2	60	30	30	30	in

Figure 11-9. Example of show port-access authenticator config Command

Table 11-2. Field Descriptions of show port-access authenticator config Command Output (Figure 11-9)

Field	Description
Port-access authenticator activated	Whether 802.1X authentication is enabled or disabled on specified port(s).
Port	Port number on switch.
Re-auth Period	Period of time (in seconds) after which clients connected to the port need to be re-authenticated.
Access Control	Port's authentication mode: Auto: Network access is allowed to any connected device that supports 802.1X authentication and provides valid 802.1X credentials. Authorized: Network access is allowed to any device connected to the port, regardless of whether it meets 802.1X criteria. Unauthorized: Network access is blocked to any device connected to the port, regardless of whether the device meets 802.1X criteria.
Max reqs	Number of authentication attempts that must time-out before authentication fails and the authentication session ends.
Quiet Period	Period of time (in seconds) during which the port does not try to acquire a supplicant.
TX Timeout	Period of time (in seconds) that the port waits to retransmit the next EAPOL PDU during an authentication session.
Suppliant Timeout	Period of time (in seconds) that the switch waits for a supplicant response to an EAP request.
Server Timeout	Period of time (in seconds) that the switch waits for a server response to an authentication request.
Cntrl Dir	Directions in which flow of incoming and outgoing traffic is blocked on 802.1X-aware port that has not yet entered the authenticated state: Both: Incoming and outgoing traffic is blocked on port until authentication occurs. In: Only incoming traffic is blocked on port before authentication occurs. Outgoing traffic with unknown destination addresses is flooded on the unauthenticated 802.1X-aware port.

Viewing 802.1X Open VLAN Mode Status

You can examine the switch's current VLAN status by using the **show port-access authenticator vlan** and **show port-access authenticator < port-list >** commands. Figure 11-12 shows related VLAN data that can help you to see how the switch is using statically configured VLANs to support 802.1X operation.

```
ProCurve(config)# show port-access authenticator vlan
```

Port Access Authenticator VLAN Configuration

Port-access authenticator activated [No] : No

Allow RADIUS-assigned dynamic (GVRP) VLANs [No] : No

Port	Access Control	Unauth VLAN ID	Auth VLAN ID
1	Auto	100	101
2	Auto	100	101
3	Auto	100	0
4	Auto	100	101

Figure 11-10. Example Showing Ports Configured for Open VLAN Mode

```
ProCurve (config)# show port-access authenticator 1-3
```

Port Access Authenticator Status

Port-access authenticator activated [No] : No

Allow RADIUS-assigned dynamic (GVRP) VLANs [No] : No

Port	Status	Current VLAN ID	Current Port COS	RADIUS ACL Applied?
1	Closed	100	No-override	No
2	Open	101	No-override	No

Figure 11-11. Example Showing Ports Configured for Open VLAN Mode

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.

Table 11-3. Output for Determining Open VLAN Mode Status (Figure 11-10)

Status Indicator	Meaning
Access Control	<p>This state is controlled by the following port-access command syntax:</p> <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>Authorized: 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>Unauthorized: 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>
Unauthorized VLAN ID	<p>< vlan-id >: Lists the VID of the static VLAN configured as the unauthorized VLAN for the indicated port.</p> <p>0: No unauthorized VLAN has been configured for the indicated port.</p>
Authorized VLAN ID	<p>< vlan-id >: Lists the VID of the static VLAN configured as the authorized VLAN for the indicated port.</p> <p>0: No authorized VLAN has been configured for the indicated port.</p>

Table 11-4. Output for Determining Open VLAN Mode Status(Figure11-11)

Status Indicator	Meaning
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>
Current Port CoS	The status of Class of Service for the port.

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.

Configuring Port-Based and User-Based Access Control (802.1X)

Displaying 802.1X Configuration, Statistics, and Counters

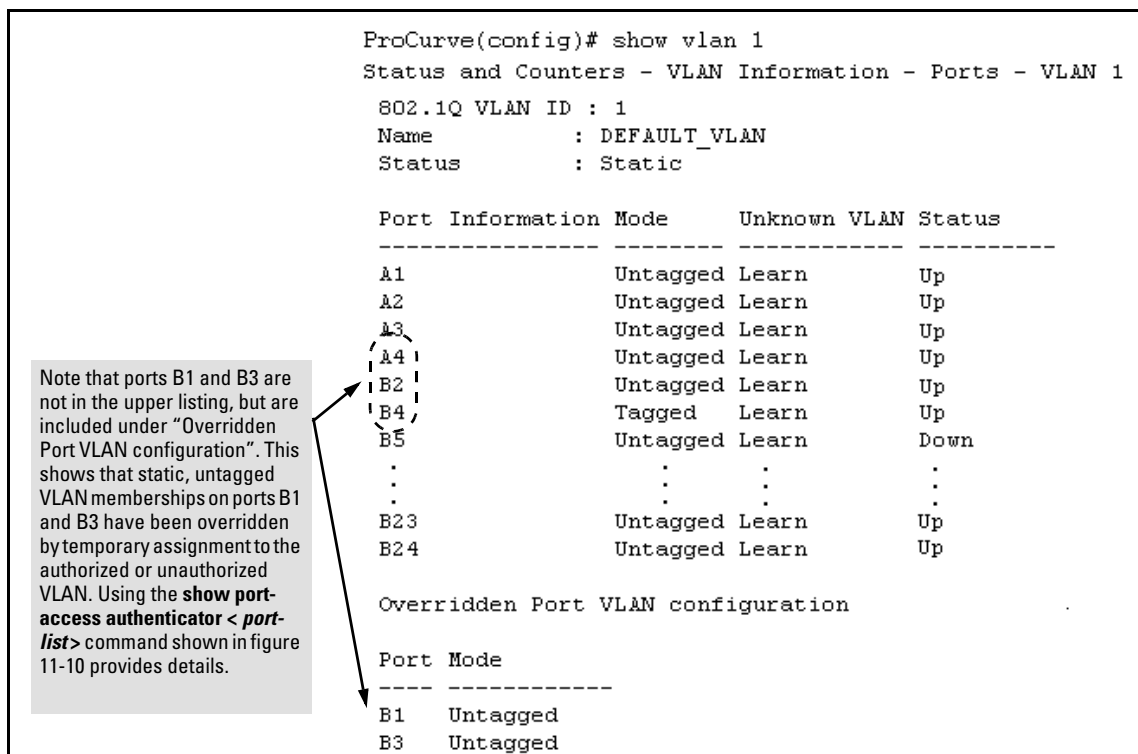


Figure 11-12. Example of Showing a VLAN with Ports Configured for Open VLAN Mode

Show Commands for Port-Access Supplicant

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

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

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

Connecting - Starting authentication.

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

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

Authenticating - Authentication is in progress.

Held - Authenticator sent notice of failure. The supplicant port is waiting for the authenticator's held-period (page 11-48).

For descriptions of the supplicant parameters, refer to "Configuring a Supplicant Switch Port" on page 11-48.

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

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

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

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

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

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

How RADIUS/802.1X Authentication Affects VLAN Operation

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

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

Using a RADIUS server to authenticate clients, you can provide port-level security protection from unauthorized network access for the following authentication methods:

- 802.1X: Port-based or client-based access control to open a port for client access after authenticating valid user credentials.
- MAC address: Authenticates a device's MAC address to grant access to the network.
- Web-browser interface: Authenticates clients for network access using a web page for user login.

Note

You can use 802.1X (port-based or client-based) authentication and either Web or MAC authentication at the same time on a port, with a maximum of 8 clients allowed on the port. (The default is one client.) Web authentication and MAC authentication are mutually exclusive on the same port. Also, you must disable LACP on ports configured for any of these authentication methods. For more information, see “Web and MAC Authentication” on page 3-1 in this guide.

VLAN Assignment on a Port

Following client authentication, VLAN configurations on a port are managed as follows when you use 802.1X, MAC, or Web authentication:

- The port resumes membership in any tagged VLANs for which it is already assigned in the switch configuration. Tagged VLAN membership allows a port to be a member of multiple VLANs simultaneously.
- The port is temporarily assigned as a member of an untagged (static or dynamic) VLAN for use during the client session according to the following order of options.
 - a. The port joins the VLAN to which it has been assigned by a RADIUS server during client authentication.
 - b. If RADIUS authentication does not include assigning the port to a VLAN, then the switch assigns the port to the authorized-client VLAN configured for the authentication method.
 - c. If the port does not have an authorized-client VLAN configured, but is configured for membership in an untagged VLAN, the switch assigns the port to this untagged VLAN.

Operating Notes

- During client authentication, a port assigned to a VLAN by a RADIUS server or an authorized-client VLAN configuration is an untagged member of the VLAN for the duration of the authenticated session. This applies even if the port is also configured in the switch as a tagged member of the same VLAN. The following restrictions apply:
 - If the port is assigned as a member of an untagged *static* VLAN, the VLAN must already be configured on the switch. If the static VLAN configuration does not exist, the authentication fails.
 - If the port is assigned as a member of an untagged *dynamic* VLAN that was learned through GVRP, the dynamic VLAN configuration must exist on the switch at the time of authentication and GVRP-learned dynamic VLANs for port-access authentication must be enabled.

If the dynamic VLAN does not exist or if you have not enabled the use of a dynamic VLAN for authentication sessions on the switch, the authentication fails.

- To enable the use of a GVRP-learned (dynamic) VLAN as the untagged VLAN used in an authentication session, enter the **aaa port-access gvrp-vlans** command, as described in “Enabling the Use of GVRP-Learned Dynamic VLANs in Authentication Sessions” on page 11-64.
- Enabling the use of dynamic VLANs in an authentication session offers the following benefits:
 - You avoid the need of having static VLANs pre-configured on the switch.
 - You can centralize the administration of user accounts (including user VLAN IDs) on a RADIUS server.

For information on how to enable the switch to dynamically create 802.1Q-compliant VLANs on links to other devices using the GARP VLAN Registration Protocol (GVRP), see “GVRP” on page 3-1 in the *Advanced Traffic Management Guide*.

- For an authentication session to proceed, a port must be an untagged member of the (static or dynamic) VLAN assigned by the RADIUS server (or an authorized-client VLAN configuration). The port temporarily drops any current untagged VLAN membership.

If the port is not already a member of the RADIUS-assigned (static or dynamic) untagged VLAN, the switch temporarily reassigns the port as an untagged member of the required VLAN (for the duration of the session). *At the same time, if the port is already configured as an untagged member of a different VLAN, the port loses access to the other VLAN for the duration of the session.* (A port can be an untagged member of only one VLAN at a time.)

When the authentication session ends, the switch removes the temporary untagged VLAN assignment and re-activates the temporarily disabled, untagged VLAN assignment.

- If GVRP is already enabled on the switch, the temporary untagged (static or dynamic) VLAN created on the port for the authentication session is advertised as an existing VLAN.

If this temporary VLAN assignment causes the switch to disable a different untagged static or dynamic VLAN configured on the port (as described in the preceding bullet and in “Example of Untagged VLAN Assignment in a RADIUS-Based Authentication Session” on page 11-61), the disabled VLAN assignment is not advertised. When the authentication session ends, the switch:

- Removes the temporary untagged VLAN assignment and stops advertising it.
 - Re-activates and resumes advertising the temporarily disabled, untagged VLAN assignment.
- If you modify a VLAN ID configuration on a port during an 802.1X, MAC, or Web authentication session, the changes do not take effect until the session ends.
- When a switch port is configured with RADIUS-based authentication to accept multiple 802.1X and/or MAC or Web authentication client sessions, all authenticated clients must use the same port-based, untagged VLAN membership assigned for the earliest, currently active client session.

Therefore, on a port where one or more authenticated client sessions are already running, all such clients are on the same untagged VLAN. If a RADIUS server subsequently authenticates a new client, but attempts to re-assign the port to a different, untagged VLAN than the one already in use for the previously existing, authenticated client sessions, the connection for the new client will fail.

Example of Untagged VLAN Assignment in a RADIUS-Based Authentication Session

The following example shows how an untagged static VLAN is temporarily assigned to a port for use during an 802.1X authentication session. In the example, an 802.1X-aware client on port A2 has been authenticated by a RADIUS server for access to VLAN 22. However, port A2 is not configured as a member of VLAN 22 but as a member of untagged VLAN 33 as shown in Figure 11-13.

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	Forbidden	Forbidden	Forbidden
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 11-13. Example of an Active VLAN Configuration

In Figure 11-13, if RADIUS authorizes an 802.1X client on port A2 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).

To view the temporary VLAN assignment as a change in the active configuration, use the **show vlan <vlan-id>** command as shown in Figure 11-12 where **<vlan-id>** is the (static or dynamic) VLAN used in the authenticated client session.

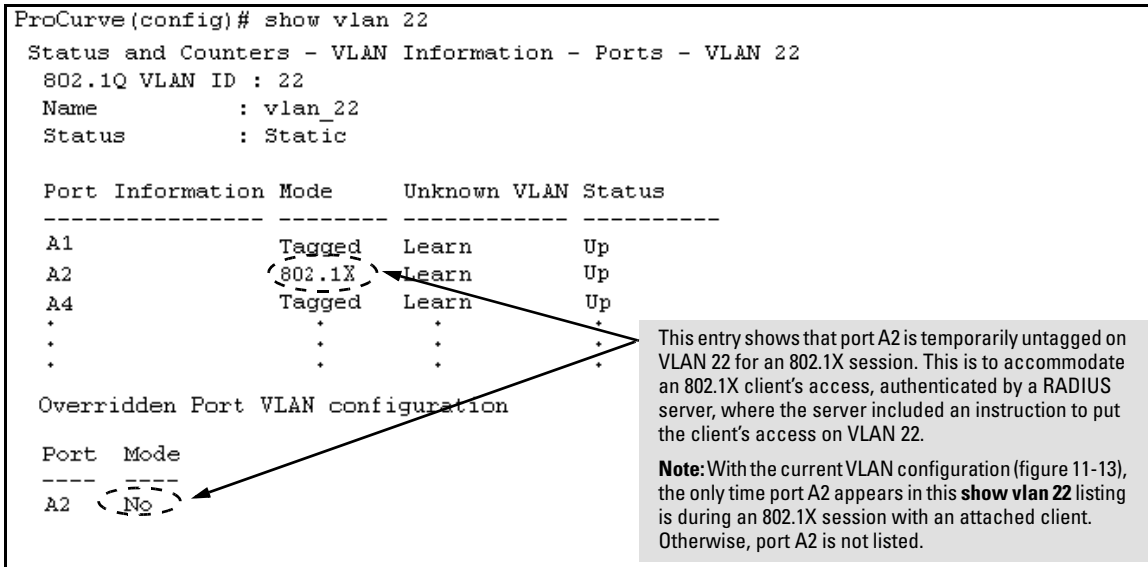


Figure 11-14. The Active Configuration for VLAN 22 Temporarily Changes for the 802.1X Session

However, as shown in Figure 11-13, because VLAN 33 is configured as untagged on port A2 and because a port can be untagged on only one VLAN, port A2 loses access to VLAN 33 for the duration of the 802.1X session on VLAN 22.

You can verify the temporary loss of access to VLAN 33 by entering the **show vlan 33** command as shown in Figure 11-15.

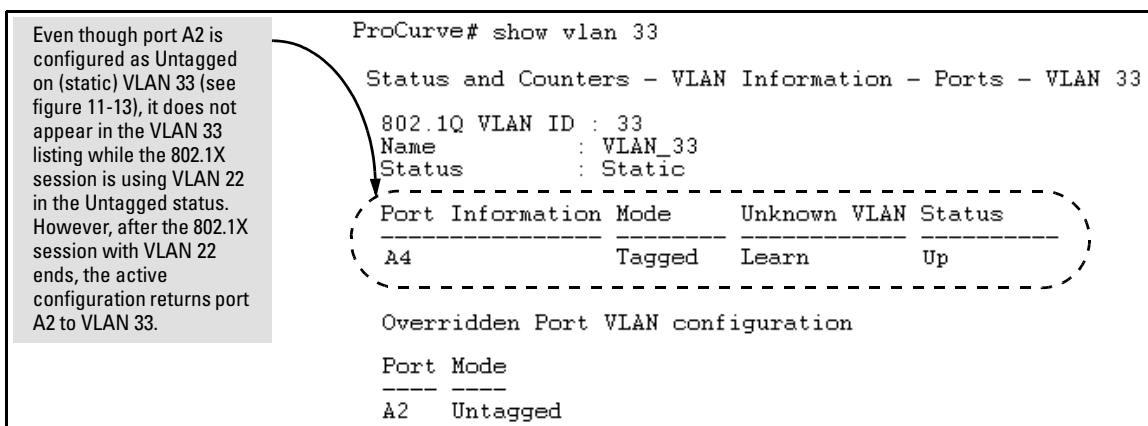


Figure 11-15. The Active Configuration for VLAN 33 Temporarily Drops Port 22 for the 802.1X Session

When the 802.1X client's session on port A2 ends, the port removes the temporary untagged VLAN membership. The static VLAN (VLAN 33) that is "permanently" configured as untagged on the port becomes available again. Therefore, 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 as shown in Figure 11-16.

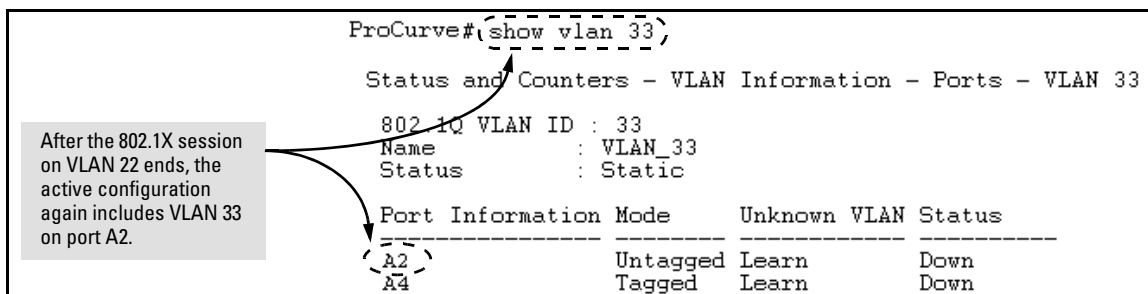


Figure 11-16.The Active Configuration for VLAN 33 Restores Port A2 After the 802.1X Session Ends

Enabling the Use of GVRP-Learned Dynamic VLANs in Authentication Sessions

Syntax: aaa port-access gvrp-vlans

Enables the use of dynamic VLANs (learned through GVRP) in the temporary untagged VLAN assigned by a RADIUS server on an authenticated port in an 802.1X, MAC, or Web authentication session.

*Enter the **no** form of this command to disable the use of GVRP-learned VLANs in an authentication session.*

For information on how to enable a switch to dynamically create 802.1Q-compliant VLANs, see "GVRP" on page 3-1 in this guide.

Notes:

1. If a port is assigned as a member of an untagged dynamic VLAN, the dynamic VLAN configuration must exist at the time of authentication and GVRP for port-access authentication must be enabled on the switch.

If the dynamic VLAN does not exist or if you have not enabled the use of a dynamic VLAN for authentication sessions on the switch, the authentication fails.

Syntax: aaa port-access gvrp-vlans
—Continued—

2. After you enable dynamic VLAN assignment in an authentication session, it is recommended that you use the **interface unknown-vlans** command on a per-port basis to prevent denial-of-service attacks. The **interface unknown-vlans** command allows you to:

- Disable the port from sending advertisements of existing GVRP-created VLANs on the switch.
- Drop all GVRP advertisements received on the port.

For more information, refer to “GVRP” on page 3-1 in the Advanced Traffic Management Guide.

3. If you disable the use of dynamic VLANs in an authentication session using the **no aaa port-access gvrp-vlans** command, client sessions that were authenticated with a dynamic VLAN continue and are not deauthenticated.

(This behavior differs from how static VLAN assignment is handled in an authentication session. If you remove the configuration of the static VLAN used to create a temporary client session, the 802.1X, MAC, or Web authenticated client is deauthenticated.)

However, if a RADIUS-configured dynamic VLAN used for an authentication session is deleted from the switch through normal GVRP operation (for example, if no GVRP advertisements for the VLAN are received on any switch port), authenticated clients using this VLAN are deauthenticated.

Note

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.
-

Operating Note

Applying Web Authentication or MAC Authentication Concurrently with Port-Based 802.1X Authentication: While 802.1X port-based access control can operate concurrently with Web Authentication or MAC Authentication, port-based access control is subordinate to Web-Auth and MAC-Auth operation. If 802.1X operates in port-based mode and MAC or Web authentication is enabled on the same port, any 802.1X authentication has no effect on the ability of a client to access the controlled port. That is, the client's access will be denied until the client authenticates through Web-Auth or MAC-Auth on the port. Note also that a client authenticating with port-based 802.1X does not open the port in the same way that it would if Web-Auth or MAC-Auth were not enabled. That is, any non-authenticating client attempting to access the port after another client authenticates with port-based 802.1X would still have to authenticate through Web-Auth or MAC-Auth.

Messages Related to 802.1X Operation

Table 11-5. 802.1X Operating Messages

Message	Meaning
Port <port-list> is not an authenticator.	The ports in the port list have not been enabled as 802.1X authenticators. Use this command to enable the ports as authenticators: <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 as a Supplicant” on page 11-48.
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 5-40).
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 Port-Based and User-Based Access Control (802.1X)
Messages Related to 802.1X Operation

Configuring and Monitoring Port Security

Contents

Overview	12-2
Basic Operation	12-2
Eavesdrop Protection	12-3
Blocking Unauthorized Traffic	12-3
Trunk Group Exclusion	12-4
Planning Port Security	12-5
Port Security Command Options and Operation	12-6
Retention of Static MAC Addresses	12-10
Displaying Current Port Security Settings	12-10
Configuring Port Security	12-12
MAC Lockdown	12-17
Differences Between MAC Lockdown and Port Security	12-19
Deploying MAC Lockdown	12-21
MAC Lockout	12-25
Port Security and MAC Lockout	12-27
Web: Displaying and Configuring Port Security Features	12-27
Reading Intrusion Alerts and Resetting Alert Flags	12-28
Notice of Security Violations	12-28
How the Intrusion Log Operates	12-29
Keeping the Intrusion Log Current by Resetting Alert Flags	12-29
Using the Event Log To Find Intrusion Alerts	12-34
Web: Checking for Intrusions, Listing Intrusion Alerts, and Resetting Alert Flags	12-35
Operating Notes for Port Security	12-35

Overview

Feature	Default	Menu	CLI	Web
Displaying Current Port Security	n/a	—	page 12-10	page 12-27
Configuring Port Security	disabled	—	page 12-12	page 12-27
Intrusion Alerts and Alert Flags	n/a	page 12-34	page 12-32	page 12-35

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.

Eavesdrop Protection. Using either the port-security command or the switch’s web browser interface to enable port security on a given port automatically enables eavesdrop prevention on that port.

General Operation for Port Security. On a per-port basis, you can configure security measures to block unauthorized devices, and to send notice of security violations. Once 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 32 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.)

Eavesdrop Protection

Configuring port security on a given switch port automatically enables eavesdrop protection for that port. This prevents use of the port to flood unicast packets addressed to MAC addresses unknown to the switch. This blocks unauthorized users from eavesdropping on traffic intended for addresses that have aged-out of the switch's address table. (Eavesdrop prevention does not affect multicast and broadcast traffic, meaning that the switch floods these two traffic types out a given port regardless of whether port security is enabled on that port.)

Blocking Unauthorized Traffic

Unless you configure the switch to disable a port on which a security violation is detected, the switch security measures block unauthorized traffic without disabling the port. This implementation enables you to apply the security

configuration to ports on which hubs, switches, or other devices are connected, and to maintain security while also maintaining network access to authorized users. For example:

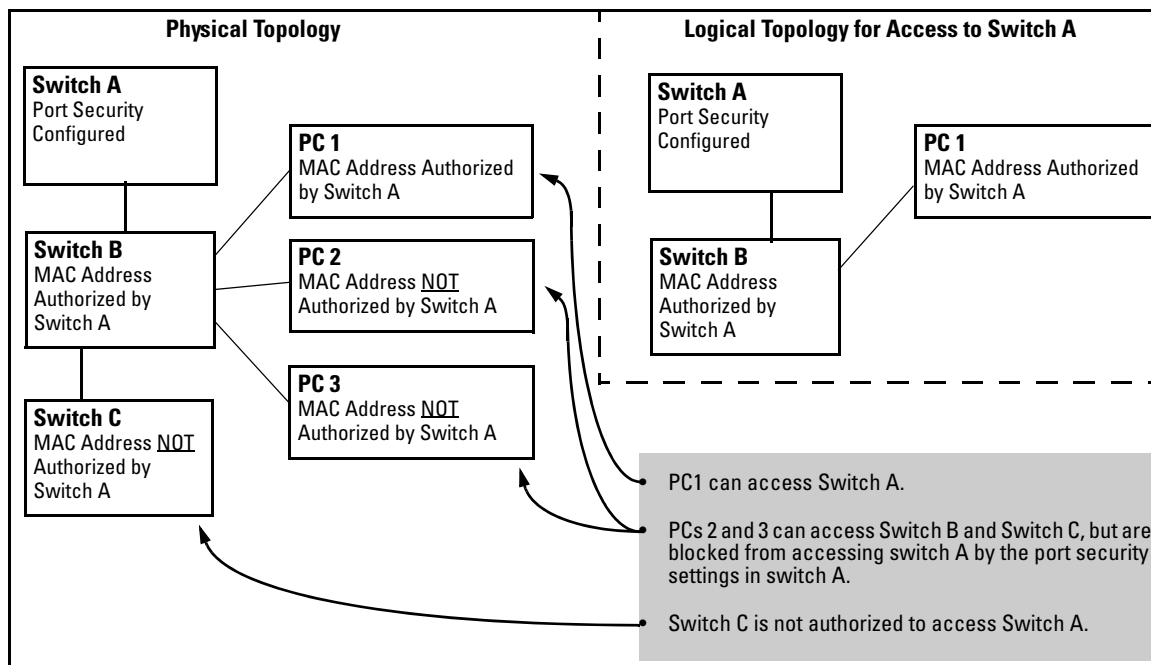


Figure 12-1. Example of How Port Security Controls Access

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 32)?
 - 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	12-11
port-security	12-12
< [ethernet] <i>port-list</i> >	12-12
[learn-mode]	12-12
[address-limit]	12-12
[mac-address]	12-12
[action]	12-12
[clear-intrusion-flag]	12-12
no port-security	12-12

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. **080090-1362f23. 080071-0c45a1**
 2. **00f031-423fc14. 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 12-10.*

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 >
(- *Continued* -)

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.*

On switches covered in this guide, automatically invokes eavesdrop protection. (Refer to “Eavesdrop Protection” on page 12-3.)

Port-Access: *Enables you to use Port Security with (802.1X) Port-Based Access Control.*

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 32.*

mac-address < mac-addr >

*Available for **static** (static-learn and configured-learn) modes. Allows up to 32 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 12-29.

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 12-28.)

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 12-2. 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 12-3. 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 12-8.

```
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 determined 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:

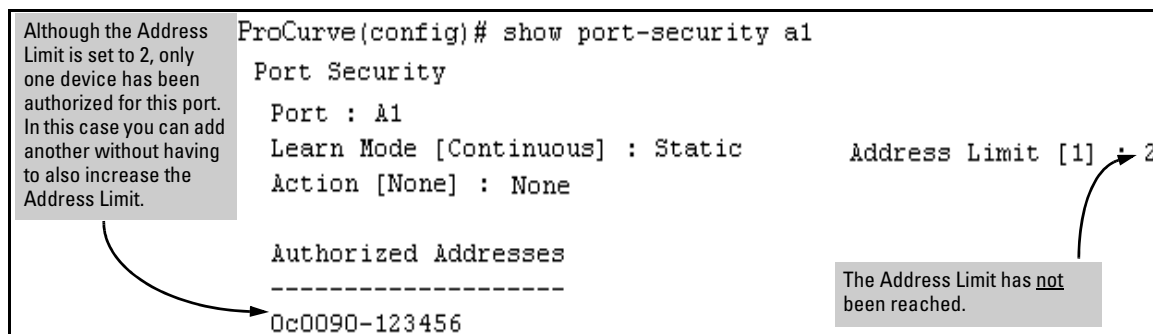


Figure 12-4. Example of Adding an Authorized Device to a Port

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:

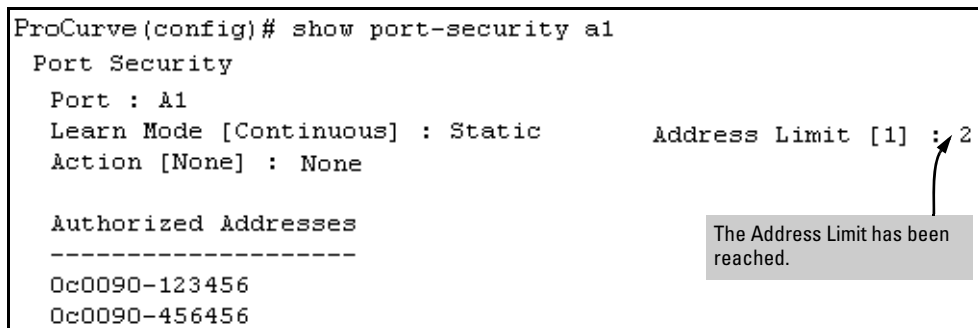


Figure 12-5. Example of Adding a Second Authorized Device to a Port

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 12-6. 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 12-8.)

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.)

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.)

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 a1
Port Security
Port : A1
Learn Mode [Continuous] : Static      Address Limit [1] : 2
Action [None] : None

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

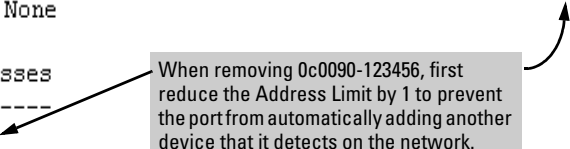


Figure 12-7. Example of Two Authorized Addresses on Port A1

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-sec a1
Port Security
  Port : A1
  Learn Mode : Static           Address Limit : 1
  Action : None
  Authorized Addresses
  -----
  0c0090-456456
```

Figure 12-8. 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.

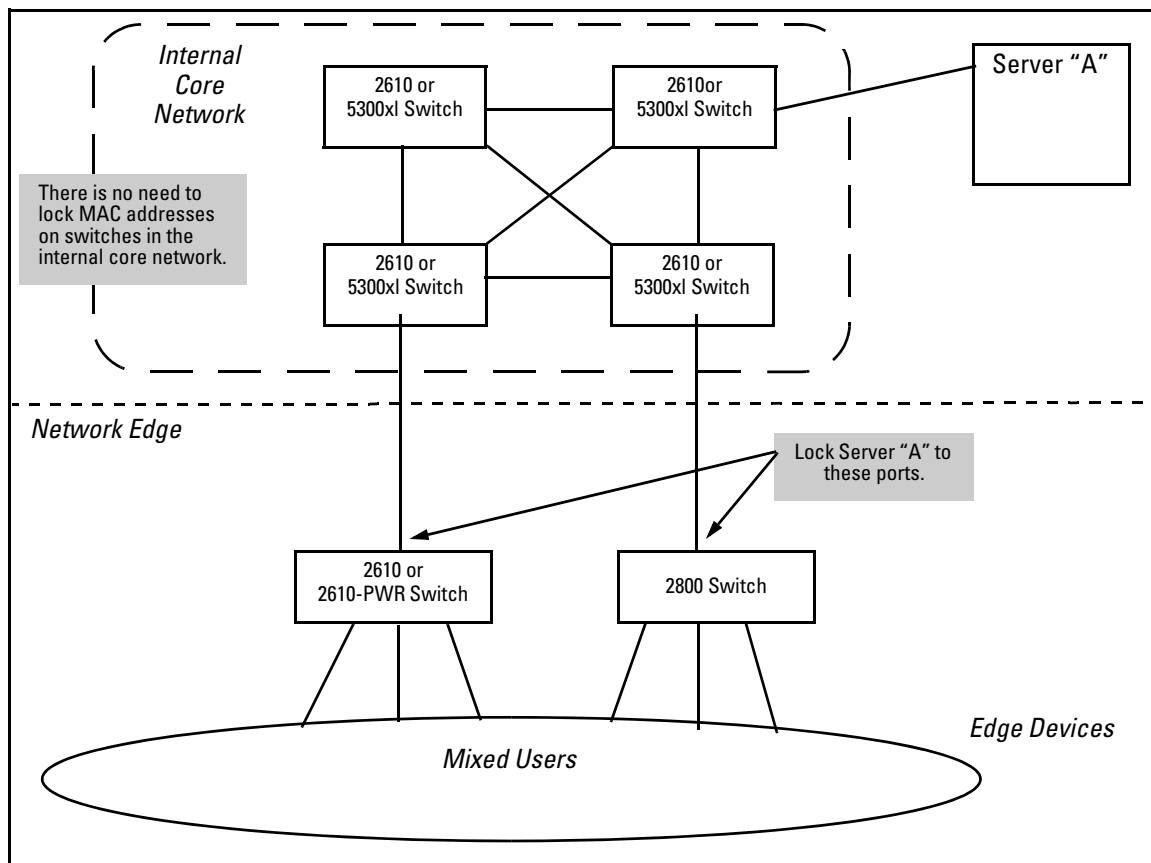


Figure 12-9. MAC Lockdown Deployed At the Network Edge Provides Security

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.

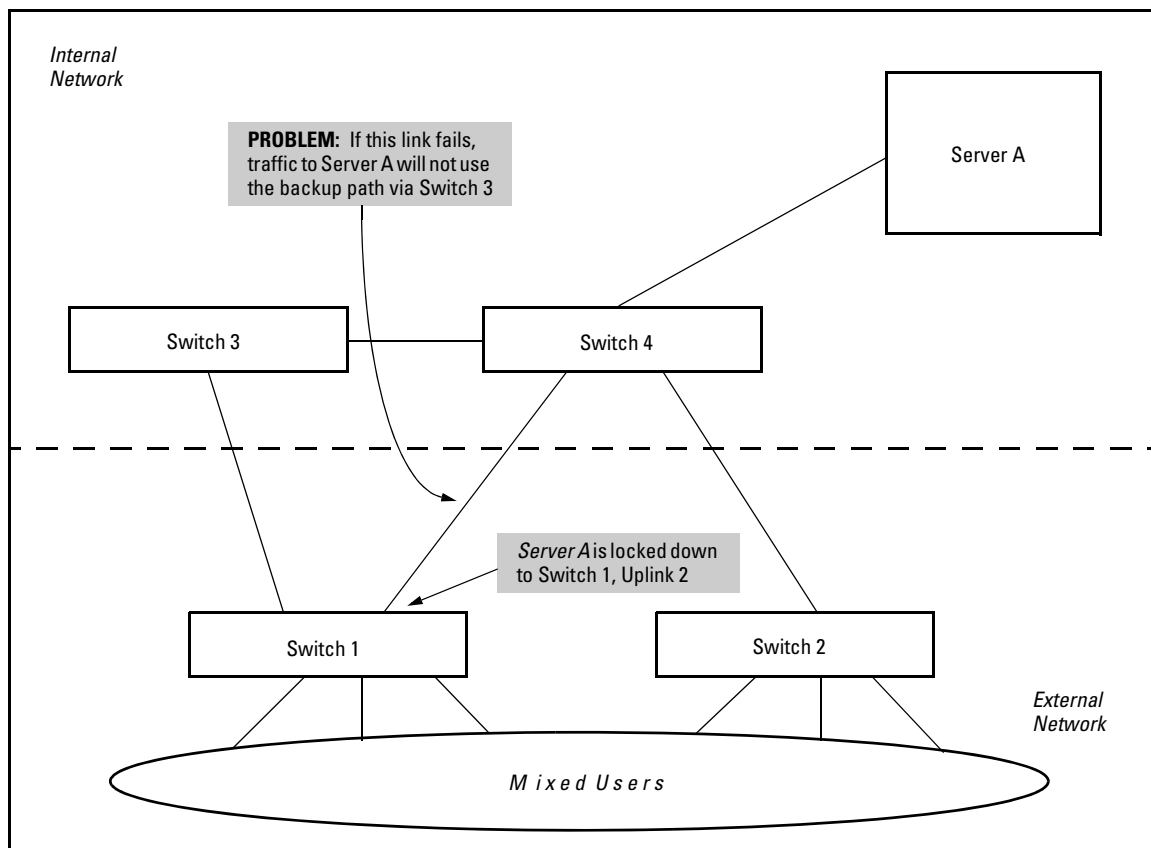


Figure 12-10. Connectivity Problems Using MAC Lockdown with Multiple Paths

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 12-9 (page 12-22), 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# show static-mac
VLAN  MAC Address Port
1 001083-34f8fa 9
Number of locked down MAC addresses = 1
```

Figure 12-11.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.

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 12-20.)

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

Figure 12-12.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 12-13. 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 12-14. Example of Port Status Screen with Intrusion Alert on Port A3

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

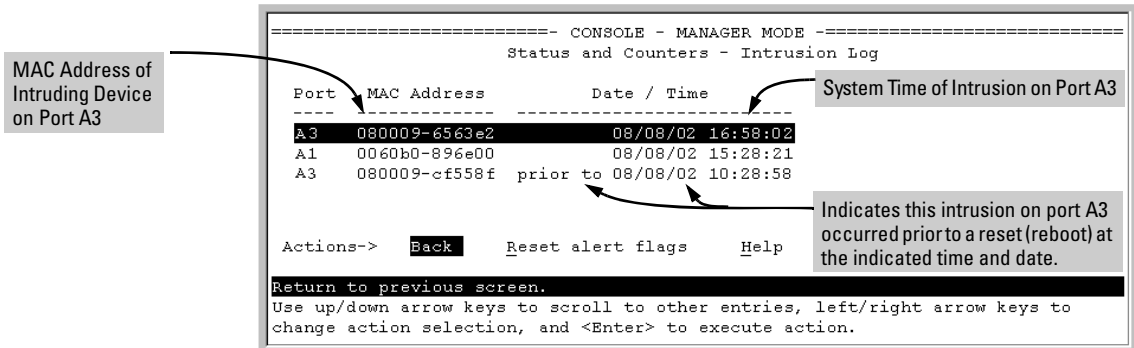


Figure 12-15. Example of the Intrusion Log Display

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 12-14 on page 12-30) 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

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 12-15, 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 12-35.)

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 Alert	Enabled	Status	Mode	Flow Ctrl
----	-----	+	-----	-----	-----	-----
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

Figure 12-16.Example of an Unacknowledged Intrusion Alert in a Port Status Display

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

MAC Address of latest Intruder on Port A1	ProCurve# show port-security intrusion-log Status and Counters - Intrusion Log	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).	<table border="1"> <thead> <tr> <th>Port</th><th>MAC Address</th><th>Date / Time</th></tr> </thead> <tbody> <tr> <td>A1</td><td>080009-e93d4f</td><td>07/03/02 21:09:34</td></tr> <tr> <td>A1</td><td>080009-21ae84</td><td>07/03/02 17:26:27</td></tr> <tr> <td>A1</td><td>080009-e93d4f prior to</td><td>07/03/02 17:18:43</td></tr> </tbody> </table>	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	
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 12-17.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
```

Intrusion Alert on port A1 is now cleared.

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

Figure 12-18.Example of Port Status Screen After Alert Flags Reset

For more on clearing intrusions, see “Note on Send-Disable Operation” on page 12-29

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
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 →

Log Command with “security” for Search String

Figure 12-19.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 13, “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.

Using Authorized IP Managers

Contents

Overview	13-2
Configuration Options	13-3
Access Levels	13-3
Defining Authorized Management Stations	13-4
Overview of IP Mask Operation	13-4
Menu: Viewing and Configuring IP Authorized Managers	13-5
CLI: Viewing and Configuring Authorized IP Managers	13-6
Web: Configuring IP Authorized Managers	13-9
Web Proxy Servers	13-9
Web-Based Help	13-10
Building IP Masks	13-10
Configuring One Station Per Authorized Manager IP Entry	13-10
Configuring Multiple Stations Per Authorized Manager IP Entry ..	13-11
Additional Examples for Authorizing Multiple Stations	13-13
Operating Notes	13-13

Overview

Authorized IP Manager Features

Feature	Default	Menu	CLI	Web
Listing (Showing) Authorized Managers	n/a	page 13-5	page 13-6	page 13-9
Configuring Authorized IP Managers	None	page 13-5	page 13-6	page 13-9
Building IP Masks	n/a	page 13-10	page 13-10	page 13-10
Operating and Troubleshooting Notes	n/a	page 13-13	page 13-13	page 13-13

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.

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

Note

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.

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 13-10.)
- **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 13-11.)

To configure the switch for authorized manager access, enter the appropriate *Authorized Manager IP* value, specify an *IP Mask*, and select either **Manager** or **Operator** for the *Access Level*. The IP Mask determines how the Authorized Manager IP value is used to allow or deny access to the switch by a 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 13-10.

Note

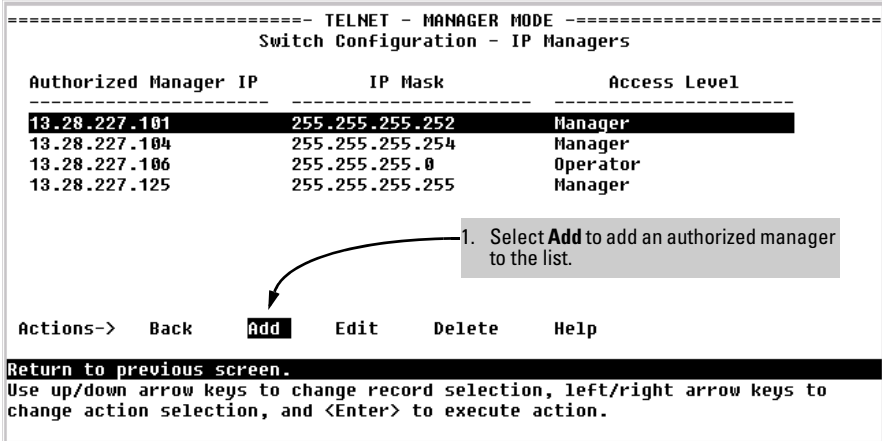
The IP Mask is a method for recognizing whether a given IP address is authorized for management access to the switch. This mask serves a different purpose than IP subnet masks and is applied in a different manner.

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 13-1. Example of How To Add an Authorized Manager Entry

Using Authorized IP Managers

Defining Authorized Management Stations

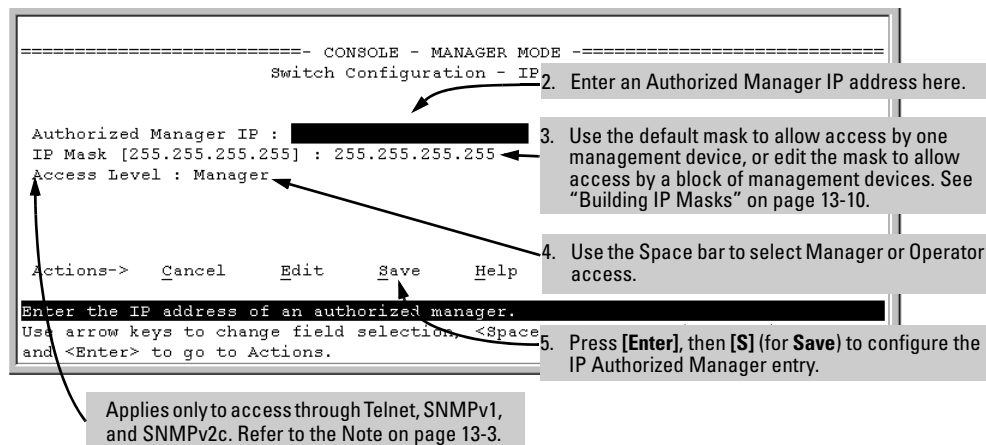


Figure 13-2. Example of How To Add an Authorized Manager Entry (Continued)

Editing or Deleting an Authorized Manager Entry. Go to the IP Managers List screen (figure 13-1), 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	13-7
<ip-address>	13-8
<ip-mask-bits>	13-8
[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 13-3. 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:




Figure 13-4. 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.

Web Proxy Servers

If you use the web browser interface to access the switch from an authorized IP manager station, it is highly recommended that you avoid using a web proxy server in the path between the station and the switch. This is because switch access through a web proxy server requires that you first add the web proxy server to the Authorized Manager IP list. This reduces security by opening switch access to anyone who uses the web proxy server.

How to Eliminate the Web Proxy Server

There are two ways to eliminate a web proxy server from the path between a station and the switch:

1. Add the IP address or DNS name of the switch to the non-proxy or Exceptions list in the web browser interface used on the authorized station.
2. If you don't need proxy server access on the authorized station, disable the proxy server feature in the station's web browser interface.

Note

IP or MAC authentication can be used without a web proxy server.

Using a Web Proxy Server to Access the Web Browser Interface

Caution

This is NOT recommended. Using a web proxy server between the stations and the switch poses a security risk. If the station uses a web proxy server to connect to the switch, any proxy user can access the switch.

If it is necessary to use the switch's web browser interface and your browser access is through a web proxy server, perform these steps:

1. Enter the web proxy server's MAC address in the port's Authorized Addresses list.
2. Enter the web proxy server's IP address in the switch's IP Authorized Managers list.

You must perform both of these steps or the switch only detects the proxy server's MAC address and IP address instead of your workstation addresses, and your connection is considered unauthorized.

Web-Based Help

For web-based help on how to use the web browser interface screen, click on the [?] button provided on the web browser screen.

Building IP Masks

The IP Mask parameter controls how the switch uses an Authorized Manager IP value to recognize the IP addresses of authorized manager stations on your network.

Configuring One Station Per Authorized Manager IP Entry

This is the easiest way to apply a mask. If you have ten or fewer management and/or operator stations, you can configure them 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 13-3 on page 13-7, 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 13-1. 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 13-2. 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	

Using Authorized IP Managers

Building IP Masks

	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 13-5, 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	






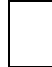
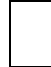

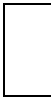





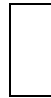

4th Octet of IP Mask:				249				
4th Octet of Authorized IP Address:				5				
Bit Numbers	Bit 7	Bit 6	Bit 5	Bit 4	Bit 3	Bit 2	Bit 1	Bit 0
Bit Values	128	64	32	16	8	4	2	1
4th Octet of IP Mask (249)								
4th Octet of IP Authorized Address (125)								
<p>Bits 1 and 2 in the mask are "off", and bits 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:</p> <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". <p>This means that stations with the IP address 13.28.227.X (where X is 121, 123, 125, or 127) are authorized.</p>								

Figure 13-5. 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.

Index

Numerics

3DES ... 7-3, 8-3

802.1X access control

- authenticate users ... 11-5

- authentication methods ... 11-4

- authentication, local ... 11-6

- authentication, user-based ... 11-4

- authenticator ... 11-18

 - operation ... 11-9

 - show commands ... 11-51

 - unblock port ... 11-6

- authorized-client VLAN, defined ... 11-6

- auth-vid ... 11-23

- auto ... 11-20

- blocked port, trunked ... 11-13

- caution, unauthorized-client VLAN ... 11-37

- CHAP ... 11-3

- chap-radius ... 11-24

- clear-statistics ... 11-26

- client, effect of disconnect ... 11-37

- client-limit, no ... 11-19

- configure

 - commands ... 11-17

 - displaying configuration ... 11-51

 - overview ... 11-16

 - port ... 11-18

- configuring method ... 11-24

- control all clients ... 11-12

- control command ... 11-20

- convert to port-based ... 11-19

- CoS override ... 11-54

- counters ... 11-51

- delay move to unauthorized-client

 - VLAN ... 11-33

- delay Unauth-Client VLAN ... 11-23

- DHCP server ... 11-38

- displaying 802.1X port configuration ... 11-53

- EAP ... 11-3

- EAPOL ... 11-7, 11-52

- eap-radius ... 11-24

- enabling controlled directions ... 11-26

- enabling on ports ... 11-18

- enabling on switch ... 11-26

- features ... 11-3

- force authorized ... 11-20, 11-55

- force unauthorized ... 11-20, 11-55

- general setup ... 11-14

- guest VLAN ... 11-7, 11-8, 11-29, 11-36

- GVRP ... 11-58, 11-59

 - effect ... 11-60, 11-65

- initialize ... 11-26

- LACP not allowed ... 11-67

- local ... 11-24

- logoff-period ... 11-23

- max-requests ... 11-21, 11-22

- MD-5 ... 11-48

- MD5 ... 11-7

- meshing, not supported ... 11-13

- messages802.1X access control

 - event log messages ... 11-67

- multiple clients ... 11-37

- multiple clients, same VLAN ... 11-5

- open port ... 11-4

- open VLAN

 - authorized client ... 11-31

 - configuration ... 11-41, 11-43

 - general operation ... 11-29

 - mode ... 11-29, 11-35

 - operating notes ... 11-44

 - operating rules ... 11-36

 - security breach ... 11-44

 - set up ... 11-40

 - status ... 11-53, 11-55

 - status, viewing ... 11-54

 - suspended VLAN ... 11-54

 - unauthorized client ... 11-31

 - use model ... 11-31

 - VLAN, after authentication ... 11-31, 11-37, 11-44

 - VLAN, tagged ... 11-31, 11-32, 11-37, 11-44, 11-56

- overview ... 11-3

- port, supplicant ... 11-15

- port-based
 - access ... 11-4
 - client without authentication ... 11-5
 - effect of Web/MAC Auth client ... 11-66
 - enable ... 11-18, 11-46
 - latest client, effect ... 11-5
 - multiple client access ... 11-6
 - multiple clients authenticating ... 11-5
 - no client limit ... 11-4
 - open port ... 11-4
 - operation ... 11-5
 - recommended use ... 11-5
 - return to ... 11-19
 - See also* user-based.
 - single client authenticates ... 11-5
 - tagged VLAN membership ... 11-5
 - unauthorized client risk ... 11-6
 - untagged VLAN membership ... 11-5, 11-30
 - with Web/MAC authentication ... 11-6
- port-security ... 11-42
- port-security use ... 11-6
- port-security, with 802.1X ... 11-45
- priority of VLAN, per-port ... 11-10, 11-30
- quiet-period ... 11-21, 11-22
- RADIUS
 - effect on VLAN operation ... 11-58, 11-59
- RADIUS ... 11-3
 - VLAN assignment ... 11-36
- RADIUS host IP address ... 11-25
- Rate-Limit override ... 11-54
- reauthenticate ... 11-26
- reauth-period ... 11-23
- rules of operation ... 11-12
- server-timeout ... 11-21, 11-22
- show commands ... 11-51
- show commands, supplicant ... 11-57
- statistics ... 11-51
- supplicant
 - client not using ... 11-34
 - configuring switch port ... 11-48
 - enabling switch port ... 11-48
 - identity option ... 11-49
 - secret ... 11-49
 - switch port operating as ... 11-47
- supplicant state ... 11-57
- supplicant statistics, note ... 11-57
- supplicant, configuring ... 11-47
- supplicant-timeout ... 11-21, 11-22

- terminology ... 11-6
- traffic flow on unauthenticated ports ... 11-27
- troubleshooting, gvrp ... 11-58, 11-59, 11-60
- trunked port blocked ... 11-13
- tx-period ... 11-21, 11-22
- unauthenticated port ... 11-26
- unauthorized ... 11-20
- unauthorized-Client VLAN ... 11-23
- unauthorized-client VLAN, defined ... 11-8
- unauthorized-Client VLAN, multiple
 - clients ... 11-39
- unauth-period ... 11-23
- unauth-period command ... 11-33
- unauth-vid ... 11-23
- use model, open VLAN mode ... 11-31
- used with port-security ... 11-45
- user-based
 - access ... 11-4

See also port based

- authentication ... 11-10
- client authentication ... 11-5
- client limit ... 11-3, 11-4, 11-46
- client-limit, enable ... 11-19
- clients use same VLAN ... 11-30
- convert to port-based ... 11-19
- enable ... 11-18, 11-46
- limit ... 11-4
- limit for web auth, MAC auth ... 11-19
- See also* user-based.
- tagged VLAN ... 11-5
- VLAN ... 11-38, 11-39
- Web/MAC Auth clients ... 11-5
- user-based vs. port-based ... 11-15
- VLAN
 - authorized-client ... 11-35, 11-36, 11-37
 - guest ... 11-36
 - RADIUS assigned, effect ... 11-38
 - RADIUS override ... 11-35
 - RADIUS-assigned ... 11-36
 - tagged ... 11-33, 11-34
 - temporary membership ... 11-36
 - unauthorized-client ... 11-36, 11-37
 - unauthorized-client, best use ... 11-39
 - unauthorized-client, caution ... 11-37
 - unauthorized-client, on different
 - ports ... 11-39
 - untagged ... 11-30, 11-33, 11-34
 - untagged membership ... 11-19

- VLAN operation ... 11-58
- VLAN use, multiple clients ... 11-6
- VLAN, assignment conflict ... 11-12
- VLAN, membership priority ... 11-10, 11-30
- VLAN, priority, RADIUS ... 11-34
- VLAN, tagged membership ... 11-34
- Wake-on-LAN traffic ... 11-27
- Web/MAC Auth effect ... 11-66

A

aaa

- authentication, authorized ... 5-9
- authorized ... 5-9
- GVRP VLANs ... 11-64

aaa authentication ... 4-8

- privilege-mode ... 4-11
- privilege-mode defined ... 4-11

aaa port-access

- See* Web or MAC Authentication.

AC

- CIDR, mask ... 9-41

access levels, authorized IP managers ... 13-3

accounting

- See* RADIUS.

ACL

- ACE sequence ... 9-38
- ACE, after match not used ... 9-36
- ACE, defined ... 9-6
- ACE, duplicates ... 9-38
- ACE, limit ... 9-24
- ACE, order in list
 - See* sequence, ACEs.
- ACL ID, defined ... 9-6
- ACL log message
 - See* ACL, logging.
- ACL mask ... 9-20
- ACL, defined ... 9-6
- application planning ... 9-16
- application, recommended ... 9-3
- applied to open connection ... 9-71
- assign to VLAN ... 9-37
- basic structure ... 9-33
- broadcasts, effect on ... 9-71
- CIDR, mask ... 9-38
- command summary ... 9-5
- command syntax ... 9-40
- configuration planning ... 9-11

- configured but not used ... 9-37
- configured, not used ... 9-37
- configuring offline ... 9-10
- copy operation appends ... 9-64
- create, CLI method ... 9-38
- DA, defined ... 9-7, 9-8
- definitions ... 9-6
- deny any, implicit ... 9-10, 9-12, 9-13, 9-15, 9-22, 9-23, 9-24, 9-33, 9-36, 9-37
- deny any, implicit, supersede ... 9-33
- deny any, implicit, switched packets ... 9-14
- deny any, rule use ... 9-17
- deny, defined ... 9-7
- editing ... 9-38
- end ... 9-37
- exit statement ... 9-37
- extended ACL, resource use ... 9-17
- extended, defined ... 9-7, 9-32
- extended, numeric I.D. range ... 9-32
- extended, structure ... 9-34
- extended, use ... 9-9
- filtering criteria ... 9-9
- filtering process ... 9-13, 9-14, 9-23
- host option ... 9-29
- i.d. range, 1-99 ... 9-39
- implicit deny
 - See* deny any, implicit.
- implicit deny, defined ... 9-7
- inbound traffic, defined ... 9-8
- logging ... 9-10, 9-11
- logging described ... 9-67
- logging, ACLs ... 9-41
- logging, performance impact ... 9-11
- logging, session ... 9-10
- managing resource use ... 9-18
- mask ... 9-10, 9-27
- mask bit overlap ... 9-20
- mask usage ... 9-16
- mask, ACL ... 9-20
- mask, CIDR ... 9-38
- mask, defined ... 9-7
- mask, multiple IP addresses ... 9-30
- mask, one IP address ... 9-29
- match, always ... 9-37
- match, criteria ... 9-28
- match, example ... 9-29
- match, ignored ... 9-23
- maximum allowed ... 9-24

- name string, maximum characters ... 9-32, 9-39
- number of entries ... 9-10
- offline creation ... 9-63
- operator, comparison ... 9-46
- outbound traffic, defined ... 9-8
- oversubscribing resources ... 9-18
- packet match, defining ... 9-20
- performance degraded ... 9-11
- permit, defined ... 9-8
- planning ... 9-11, 9-16
- policies ... 9-16
- policy application points ... 9-4
- prioritizing feature usage ... 9-16
- purpose ... 9-4
- recommended use ... 9-3
- replacing ... 9-24
- resource usage ... 9-16
- resource usage, help display ... 9-18
- resource use, example ... 9-19
- resource use, troubleshooting ... 9-18
- resource, display current use ... 9-18
- routed traffic ... 9-25
- rule and mask usage ... 9-16
- rules, configuration ... 9-24
- rules, operation ... 9-24
- SA, defined ... 9-8
- security use ... 9-4, 9-22
- security use, caution ... 9-23
- sequence, ACEs ... 9-38
- source routing, caution ... 9-11, 9-32
- standard ACL, resource use ... 9-17
- standard, defined ... 9-8, 9-32
- standard, example ... 9-42
- standard, structure ... 9-34
- standard, use ... 9-9, 9-39
- static VLAN requirement ... 9-11, 9-24, 9-25
- supernetting ... 9-27
- supersede implicit deny any ... 9-36
- switched packets ... 9-14
- syntax
 - See* command syntax.
- Syslog
 - See* ACL logging.
- TCP or UDP port number, IANA ... 9-47
- terms ... 9-6
- traffic types filtered ... 9-4, 9-11
- types, defined ... 9-32
- VLAN assignment ... 9-12

- VLANs ... 9-24
- where applied to traffic ... 9-12, 9-25
- wildcard ... 9-28, 9-29
- wildcard, defined ... 9-8
- ACL, standard numeric I.D. range ... 9-32
- ACLs
 - management access protection ... 1-3
 - See also* RADIUS-assigned ACLs.
- address
 - authorized for port security ... 12-3
- authentication
 - See* TACACS.
- authentication, RADIUS override ... 6-3
- authorized addresses
 - for IP management security ... 13-4
 - for port security ... 12-3
- authorized IP managers
 - access levels ... 13-3
 - building IP masks ... 13-10
 - configuring in browser interface ... 13-7, 13-9
 - configuring in console ... 13-5
 - definitions of single and multiple ... 13-4
 - effect of duplicate IP addresses ... 13-13
 - IP mask for multiple stations ... 13-11
 - IP mask for single station ... 13-10
 - IP mask operation ... 13-4
 - operating notes ... 13-13
 - overview ... 13-1
 - precedence over other security ... 13-2
 - troubleshooting ... 13-13
- authorized, authentication ... 5-9

C

- certificate
 - CA-signed ... 8-4
 - root ... 8-4
 - self-signed ... 8-4
- Class of Service ... 6-3, 6-4, 6-5
- RADIUS ... 6-3
- Clear button
 - to delete password protection ... 2-5
- configuration
 - port security ... 12-5
 - RADIUS
 - See* RADIUS.
 - SSH
 - See* SSH.

- connection inactivity time ... 2-3
- console, for configuring
 - authorized IP managers ... 13-5
- CoS ... 6-3, 6-4, 6-5
 - RADIUS override ... 6-4
- CoS override ... 11-54

D

- DA, defined ... 6-7, 9-8
- DES ... 7-3, 8-3
- disclaimer ... 1-ii
- duplicate IP address
 - effect on authorized IP managers ... 13-13

E

- Eavesdrop Protection ... 12-2
- enhancing network security ... 6-9
- event log
 - intrusion alerts ... 12-34

F

- filter, source-port
 - configuring ... 10-5
 - editing ... 10-9
 - filter indexing ... 10-9
 - filter type ... 10-8, 10-12
 - idx ... 10-8, 10-9, 10-12
 - index ... 10-8, 10-9, 10-12
 - multinetted VLAN ... 10-3
 - named source-port filters ... 10-10
 - operating rules ... 10-4, 10-10
 - port-trunk operation ... 10-2, 10-6
 - show ... 10-7, 10-12
 - value ... 10-8, 10-12
 - viewing ... 10-7, 10-12

G

- guest VLAN ... 11-7, 11-8, 11-29
- GVRP ... 11-58, 11-59
- GVRP VLANs ... 11-64
- GVRP, static VLAN not advertised ... 11-60, 11-65
- gvrp-vlans ... 11-64

I

- IANA ... 9-47
- Identity Driven Manager
 - See* IDM.
- IDM ... 6-2, 6-6, 6-25, 9-4
 - See also* RADIUS-assigned ACLs
 - RADIUS-assigned ACLs.
- inconsistent value, message ... 12-14
- intrusion alarms
 - entries dropped from log ... 12-35
 - event log ... 12-34
 - prior to ... 12-35
- Intrusion Log
 - prior to ... 12-31, 12-33
- IP
 - authorized IP managers ... 13-1
 - reserved port numbers ... 7-17
- IP masks
 - building ... 13-10
 - for multiple authorized manager stations ... 13-11
 - for single authorized manager station ... 13-10
 - operation ... 13-4

K

- kill command ... 7-11

L

- LACP
 - 802.1X not allowed ... 11-13, 11-18, 11-67

M

- MAC auth
 - port access ... 11-4
- MAC Authentication
 - authenticator operation ... 3-5
 - blocked traffic ... 3-4
- CHAP
 - defined ... 3-9
 - usage ... 3-4
- client status ... 3-33
- configuration commands ... 3-24

- configuring
 - on the switch ... 3-23
 - switch for RADIUS access ... 3-15
 - the RADIUS server ... 3-13
- features ... 3-4
- general setup ... 3-12
- LACP not allowed ... 3-12
- rules of operation ... 3-10
- terminology ... 3-9
- manager password ... 2-2, 2-4
- manager password recommended ... 4-7
- MD5
 - See *RADIUS*.
- message
 - inconsistent value ... 12-14

N

- NAS ... 6-7
- network management applications ... 6-2

O

- open VLAN mode
 - See 802.1X access control.
- OpenSSH ... 7-3
- OpenSSL ... 8-2
- operating notes
 - authorized IP managers ... 13-13
 - port security ... 12-35
- operator password ... 2-2, 2-4

P

- password
 - authorized IP managers, precedence ... 13-2
 - 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 ... 7-18
- PCM
 - See ProCurve Manager.
- port
 - security configuration ... 12-2
- port access
 - client limit ... 11-19
 - concurrent ... 11-19
 - MAC auth ... 11-4
 - See also 802.1X access control.
 - Web auth ... 11-4
 - Web/MAC ... 11-19
- port security
 - authorized address definition ... 12-3
 - authorized IP managers, precedence ... 13-2
 - basic operation ... 12-2
 - configuring ... 12-5
 - configuring in browser interface ... 12-27, 12-35
 - event log ... 12-34
 - notice of security violations ... 12-28
 - operating notes ... 12-35
 - overview ... 12-2
 - prior to ... 12-35
 - proxy web server ... 12-35
- port-based access control
 - authorized IP managers, precedence ... 13-2
 - See 802.1X access control.
 - VLAN, tagged member ... 11-31
- prior to ... 12-31, 12-33, 12-35
- Privacy Enhanced Mode (PEM)
 - See *SSH*.
- privilege-mode ... 4-11
- ProCurve Manager ... 6-2
- proxy
 - web server ... 12-35

Q

- quick start ... 1-8

R

- RADIUS
 - accounting ... 5-2, 5-26
 - accounting, configuration outline ... 5-28
 - accounting, configure server access ... 5-28
 - accounting, configure types on switch ... 5-30
 - accounting, exec ... 5-27, 5-30

- accounting, interim updating ... 5-32
- accounting, network ... 5-30
- accounting, operating rules ... 5-27
- accounting, server failure ... 5-27
- accounting, session-blocking ... 5-32
- accounting, start-stop method ... 5-30
- accounting, statistics terms ... 5-34
- accounting, stop-only method ... 5-31
- accounting, system ... 5-27, 5-30
- ACL, dynamic port ... 6-13
- authentication options ... 5-2
- authentication, local ... 5-17
- authorization ... 5-18
- authorized IP managers, precedence ... 13-2
- bypass RADIUS server ... 5-10
- Class of Service ... 6-3, 6-4, 6-5
- commands authorization ... 5-18
- commands, accounting ... 5-26
- commands, switch ... 5-6
- configuration outline ... 5-7
- configure server access ... 5-11
- configuring server ... 5-20
- configuring switch global parameters ... 5-13
- CoS override ... 6-3
- dynamic port ACL ... 6-7, 6-9
- general setup ... 5-5
- HP-Command-Exception ... 5-21
- HP-command-string ... 5-20
- local authentication ... 5-10
- MD5 ... 5-4
- messages ... 5-39
- multiple ACL application types in use ... 6-13
- network accounting ... 5-26
- operating rules, switch ... 5-4
- override CoS ... 6-4
- override CoS, example ... 6-4, 6-5
- override Rate-Limiting ... 6-4
- override Rate-Limiting, example ... 6-4, 6-5
- override, precedence, multiple clients ... 6-5
- rate-limiting ... 6-3
- Rate-Limiting override ... 6-3
- security ... 5-10
- server access order ... 5-27
- server access order, changing ... 5-37
- servers, multiple ... 5-14
- show accounting ... 5-36
- show authentication ... 5-35
- statistics, viewing ... 5-33
- terminology ... 5-3
- TLS ... 5-4
- vendor specific attributes ... 5-20
- vendor-specific attributes ... 6-3
- VSAs ... 5-22
- Web browser authentication ... 5-7
- web-browser access controls ... 5-18
- web-browser security not supported ... 5-18
- RADIUS accounting
 - See *RADIUS*.
- RADIUS-assigned ACLs ... 6-6
 - 802.1X port-based access ... 6-15
 - 802.1X user-based access ... 6-15
 - ACE, defined ... 6-6
 - application type ... 6-7
 - contrasting dynamic and static ... 6-11
 - DA, defined ... 6-7
 - defined ... 6-6
 - definitions ... 6-6
 - deny any, implicit, switched packets ... 6-14
 - deny, defined ... 6-7
 - dynamic port ... 6-9, 6-13
 - dynamic port ACL ... 6-7
 - dynamic port ACL, effect ... 6-15
 - filters ... 6-9
 - implicit deny, defined ... 6-7
 - inbound traffic, defined ... 6-7
 - inverse mask
 - See wildcard.
 - mask ... 6-7
 - mask, defined ... 6-7
 - multiple application types in use ... 6-13
 - multiple clients, access restriction ... 6-15
 - multiple dynamic ACLs ... 6-15
 - outbound traffic, defined ... 6-8
 - permit, defined ... 6-8
 - RADIUS-based ... 6-13
 - resource monitor ... 6-25
 - See also ACLs.
 - source routing, caution ... 6-12
 - static-port ACL ... 6-8
 - switched packets ... 6-14
 - terminology ... 6-6
 - terms ... 6-6
 - wildcard ... 6-7, 6-8
 - wildcard, defined ... 6-8
- RADIUS-based ACL filtering ... 6-13
- Rate-Limit override ... 11-54

- rate-limiting ... 6-3
- Rate-Limiting, RADIUS override ... 6-4
- reserved port numbers ... 7-17, 8-20
- routing
 - source-routing, caution ... 6-12, 9-11, 9-32

S

- security
 - authorized IP managers ... 13-1
 - per port ... 12-2
- security violations
 - notices of ... 12-28
- security, ACL
 - See* ACL, security use.
- security, password
 - See* SSH.
- setting a password ... 2-4
- setup screen ... 1-8
- show
 - locked down MAC addresses ... 12-25
 - locked out MAC addresses ... 12-26
- single sign-on ... 4-11
- source-routing, caution ... 6-12, 9-11, 9-32
- spanning tree
 - edge port configuration ... 11-27
- SSH
 - authenticating switch to client ... 7-3
 - authentication, client public key ... 7-2
 - authentication, user password ... 7-2
 - caution, restricting access ... 7-20
 - caution, security ... 7-18
 - CLI commands ... 7-9
 - client behavior ... 7-15, 7-16
 - client public-key authentication ... 7-19, 7-22
 - client public-key, clearing ... 7-26
 - client public-key, creating file ... 7-23
 - client public-key, displaying ... 7-25
 - configuring authentication ... 7-18
 - crypto key ... 7-11
 - disabling ... 7-11
 - enable ... 7-16, 8-19
 - enabling ... 7-15
 - erase host key pair ... 7-11
 - generate host key pair ... 7-11
 - generating key pairs ... 7-10
 - host key pair ... 7-11
 - key, babble ... 7-11

- key, fingerprint ... 7-11
- keys, zeroing ... 7-11
- key-size ... 7-17
- known-host file ... 7-13, 7-15
- man-in-the-middle spoofing ... 7-16
- messages, operating ... 7-28
- OpenSSH ... 7-3
- operating rules ... 7-8
- outbound SSH not secure ... 7-8
- password security ... 7-18
- password-only authentication ... 7-18
- passwords, assigning ... 7-9
- PEM ... 7-4
- prerequisites ... 7-5
- public key ... 7-5, 7-13
- public key, displaying ... 7-14
- reserved IP port numbers ... 7-17
- security ... 7-18
- SSHv1 ... 7-2
- SSHv2 ... 7-2
- stacking, security ... 7-8
- steps for configuring ... 7-6
- supported encryption methods ... 7-3
- switch key to client ... 7-12
- terminology ... 7-4
- unauthorized access ... 7-20, 7-27
- version ... 7-2
- zeroing a key ... 7-11
- zeroize ... 7-11

SSL

- CA-signed ... 8-4, 8-15
- CA-signed certificate ... 8-4, 8-15
- CLI commands ... 8-7
- client behavior ... 8-17, 8-18
- crypto key ... 8-10
- disabling ... 8-10
- enabling ... 8-17
- erase certificate key pair ... 8-10
- erase host key pair ... 8-10
- generate CA-signed certificate ... 8-15
- generate host key pair ... 8-10
- generate self-signed ... 8-13
- generate self-signed certificate ... 8-10, 8-13
- generate server host certificate ... 8-10
- generating Host Certificate ... 8-9
- host key pair ... 8-10
- key, babble ... 8-12
- key, fingerprint ... 8-12

- man-in-the-middle spoofing ... 8-18
- OpenSSL ... 8-2
- operating notes ... 8-6
- operating rules ... 8-6
- passwords, assigning ... 8-7
- prerequisites ... 8-5
- remove self-signed certificate ... 8-10
- remove server host certificate ... 8-10
- reserved TCP port numbers ... 8-20
- root ... 8-4
- root certificate ... 8-4
- self-signed ... 8-4, 8-13
- self-signed certificate ... 8-4, 8-10, 8-13
- server host certificate ... 8-10
- SSL server ... 8-3
- SSLv3 ... 8-2
- stacking, security ... 8-6
- steps for configuring ... 8-5
- supported encryption methods ... 8-3
- terminology ... 8-3
- TLSv1 ... 8-2
- troubleshooting, operating ... 8-21
- version ... 8-2
- zeroize ... 8-10, 8-12
- stacking
 - SSH security ... 7-8
 - SSL security ... 8-6
- STP
 - prerequisite for 802.1X controlled directions ... 11-27
- supernetting ... 9-27
- supersede implicit deny any ... 9-33
- Syslog
 - See* ACL, logging.

T

TACACS

- aaa parameters ... 4-12
- authentication ... 4-3
- authentication process ... 4-22
- authentication, local ... 4-24
- authorized IP managers, effect ... 4-28
- authorized IP managers, precedence ... 13-2
- configuration, authentication ... 4-10
- configuration, encryption key ... 4-21
- configuration, server access ... 4-17
- configuration, timeout ... 4-22

- configuration, viewing ... 4-10
- encryption key ... 4-6, 4-17, 4-18, 4-21
- encryption key, general operation ... 4-25
- encryption key, global ... 4-22
- general operation ... 4-2
- IP address, server ... 4-17
- local manager password requirement ... 4-28
- messages ... 4-27
- NAS ... 4-3
- overview ... 1-2
- precautions ... 4-5
- preparing to configure ... 4-8
- preventing switch lockout ... 4-17
- privilege level code ... 4-7
- server access ... 4-17
- server priority ... 4-20
- setup, general ... 4-5
- show authentication ... 4-8
- single login ... 4-13
- single sign-on ... 4-13
- system requirements ... 4-5
- TACACS+ server ... 4-3
- testing ... 4-5
- timeout ... 4-17
- troubleshooting ... 4-6
- unauthorized access, preventing ... 4-7
- web access, controlling ... 4-26
- web access, no effect on ... 4-5

tacacs-server ... 4-8

TCP

- reserved port numbers ... 8-20

TLS

See RADIUS.

troubleshooting

- authorized IP managers ... 13-13

trunk

- filter, source-port ... 10-2, 10-6
- LACP, 802.1X not allowed ... 11-18
- See also* LACP.

U

user name

- cleared ... 2-5

V

- value, inconsistent ... 12-14

- vendor specific attributes ... 5-22
- Vendor-Specific Attribute ... 6-8
- vendor-specific attribute
 - configuring ... 6-3
- vendor-specific attributes ... 6-3
- VLAN
 - 802.1X ... 11-58
 - 802.1X, ID changes ... 11-61, 11-65
 - 802.1X, suspend untagged VLAN ... 11-54
 - filter, source-port ... 10-3
 - not advertised for GVRP ... 11-60, 11-65
- VSA ... 6-8
 - See* vendor-specific attribute.
- VSAs, defining ... 5-22

W

- Wake-on-LAN
 - on 802.1X-aware ports ... 11-27
- warranty ... 1-ii
- Web auth
 - port access ... 11-4
- Web Authentication
 - authenticator operation ... 3-5
 - blocked traffic ... 3-4
 - CHAP
 - defined ... 3-9
 - usage ... 3-4
 - client status ... 3-33
 - configuration commands ... 3-19
 - configuring
 - on the switch ... 3-18
 - switch for RADIUS access ... 3-15
 - features ... 3-4
 - general setup ... 3-12
 - LACP not allowed ... 3-12
 - redirect URL ... 3-9
 - rules of operation ... 3-10
 - show status and configuration ... 3-28
 - terminology ... 3-9
- web browser interface, for configuring
 - authorized IP managers ... 13-7, 13-9
- web browser interface, for configuring port
 - security ... 12-27, 12-35
- web server, proxy ... 12-35
- wildcard
 - See* ACL.
- wildcard, ACL, defined ... 6-8, 9-8



© Copyright 2007 Hewlett-Packard
Development Company, L.P.

December 2007

Manual Part Number
5991-8642