



6200yl
5400zl
3500yl

Access Security Guide

ProCurve Switches
K.12.XX

www.procurve.com



ProCurve

Series 5400zl Switches

Series 3500yl Switches

6200yl Switch

February 2007

K.12.XX

Access Security Guide

© Copyright 2005-2007 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. All Rights Reserved.

This document contains proprietary information, which is protected by copyright. No part of this document may be photocopied, reproduced, or translated into another language without the prior written consent of Hewlett-Packard.

Publication Number

5991-3828
February 2007

Applicable Products

ProCurve Switch 5406zl	(J8697A)
ProCurve Switch 5412zl	(J8698A)
ProCurve Switch 3500yl-24G-PWR Intelligent Edge	(J8692A)
ProCurve Switch 3500yl-48G-PWR Intelligent Edge	(J8693A)
ProCurve Switch 6200yl-24G	(J8992A)

Trademark Credits

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

Software Credits and Notices

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 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 www.openssl.org.

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

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

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

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

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

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

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

Disclaimer

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

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

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

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

Warranty

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

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

Contents

Product Documentation

About Your Switch Manual Set	xix
Printed Publications.....	xix
Electronic Publications.....	xix
Software Feature Index	xx

1 Security Overview

Contents	1-1
Introduction	1-2
About This Guide	1-2
For More Information	1-2
Switch Access Security	1-3
Default Configuration Settings and Access Security	1-3
Local Manager Password	1-3
Inbound Telnet Access and Web Browser Access	1-4
SNMP Access (Simple Network Management Protocol)	1-4
Front-Panel Access and Physical Security	1-6
Secure File Transfers	1-6
Other Provisions for Management Access Security	1-7
Authorized IP Managers	1-7
Secure Management VLAN	1-7
TACACS+ Authentication	1-7
RADIUS Authentication	1-7
ACLs for Management Access Protection	1-7
Network Security Features	1-8
Access Control Lists (ACLs)	1-8
802.1X Access Control	1-8
Web and MAC Authentication	1-9
Secure Shell (SSH)	1-9
Secure Socket Layer (SSLv3/TLSv1)	1-10

Traffic/Security Filters	1-10
Port Security, MAC Lockdown, and MAC Lockout	1-10
Key Management System (KMS)	1-11
Advanced Threat Detection	1-12
BPDU Filtering and BPDU Protection	1-12
Connection-Rate Filtering Based On Virus-Throttling Technology	1-12
DHCP Snooping, Dynamic ARP Protection, and Instrumentation Monitor	1-12
Identity-Driven Manager (IDM)	1-13

2 Configuring Username and Password Security

Contents	2-1
Overview	2-2
Configuring Local Password Security	2-5
Menu: Setting Passwords	2-5
CLI: Setting Passwords and Usernames	2-7
Web: Setting Passwords and Usernames	2-8
SNMP: Setting Passwords and Usernames	2-8
Front-Panel Security	2-8
When Security Is Important	2-9
Front-Panel Button Functions	2-10
Clear Button	2-10
Reset Button	2-11
Restoring the Factory Default Configuration	2-11
Configuring Front-Panel Security	2-12
Disabling the Clear Password Function of the Clear Button on the Switch's Front Panel	2-14
Re-Enabling the Clear Button on the Switch's Front Panel and Setting or Changing the "Reset-On-Clear" Operation	2-16
Changing the Operation of the Reset+Clear Combination	2-17
Password Recovery	2-18
Disabling or Re-Enabling the Password Recovery Process	2-18
Password Recovery Process	2-20

3 Virus Throttling

Contents	3-1
Overview of Connection-Rate Filtering	3-3
Features and Benefits	3-4
General Operation	3-5
Filtering Options	3-5
Sensitivity to Connection Rate Detection	3-5
Application Options	3-6
Operating Rules	3-7
Unblocking a Currently Blocked Host	3-7
General Configuration Guidelines	3-8
For a network that is relatively attack-free:	3-8
For a network that appears to be under significant attack:	3-9
Configuring Connection-Rate Filtering	3-10
Global and Per-Port Configuration	3-10
Enabling Connection-Rate Filtering and Configuring Sensitivity ...	3-11
Configuring the Per-Port Filtering Mode	3-12
Example of a Basic Connection-Rate Filtering Configuration ..	3-13
Viewing and Managing Connection-Rate Status	3-15
Viewing Connection-Rate Configuration	3-15
Listing Currently-Blocked Hosts	3-17
Unblocking Currently-Blocked Hosts	3-18
Configuring and Applying Connection-Rate ACLs	3-19
Connection-Rate ACL Operation	3-20
Configuring a Connection-Rate ACL Using	
Source IP Address Criteria	3-21
Configuring a Connection-Rate ACL Using UDP/TCP Criteria	3-23
Applying Connection-Rate ACLs	3-26
Using CIDR Notation To Enter the ACE Mask	3-26
Example of Using an ACL in a Connection-Rate Configuration ...	3-27
Connection-Rate ACL Operating Notes	3-29
Connection-Rate Log and Trap Messages	3-31

4 Web and MAC Authentication

Contents	4-1
Overview	4-2
Client Options	4-3
General Features	4-3
How Web and MAC Authentication Operate	4-5
Authenticator Operation	4-5
Web-based Authentication	4-5
MAC-based Authentication	4-7
Terminology	4-9
Operating Rules and Notes	4-10
General Setup Procedure for Web/MAC Authentication	4-12
Do These Steps Before You Configure Web/MAC Authentication ..	4-12
Additional Information for Configuring the RADIUS Server To Support MAC Authentication	4-13
Configuring the Switch To Access a RADIUS Server	4-14
Configuring Web Authentication on the Switch	4-17
Overview	4-17
Configure the Switch for Web-Based Authentication	4-18
Configuring MAC Authentication on the Switch	4-24
Overview	4-24
Configure the Switch for MAC-Based Authentication	4-25
Show Commands for Web-Based Authentication	4-28
Example: Verifying a Web Authentication Configuration	4-29
Configuring MAC Authentication	4-31
Configuration Overview	4-31
Config Commands for MAC-Based Authentication	4-31
Show Commands for MAC-Based Authentication	4-36
Example: Verifying a MAC Authentication Configuration	4-38
Client Status	4-39

5 TACACS+ Authentication

Contents	5-1
-----------------------	-----

Overview	5-2
Terminology Used in TACACS Applications:	5-3
General System Requirements	5-5
General Authentication Setup Procedure	5-5
Configuring TACACS+ on the Switch	5-8
Before You Begin	5-8
CLI Commands Described in this Section	5-9
Viewing the Switch's Current Authentication Configuration	5-9
Viewing the Switch's Current TACACS+ Server Contact Configuration	5-10
Configuring the Switch's Authentication Methods	5-11
Configuring the Switch's TACACS+ Server Access	5-15
How Authentication Operates	5-20
General Authentication Process Using a TACACS+ Server	5-20
Local Authentication Process	5-22
Using the Encryption Key	5-23
General Operation	5-23
Encryption Options in the Switch	5-23
Controlling Web Browser Interface	
Access When Using TACACS+ Authentication	5-24
Messages Related to TACACS+ Operation	5-25
Operating Notes	5-25
6 RADIUS Authentication and Accounting	
Contents	6-1
Overview	6-3
Authentication Services	6-3
Accounting Services	6-4
RADIUS-Administered CoS and Rate-Limiting	6-4
SNMP Access to the Switch's Authentication Configuration MIB ...	6-4
Terminology	6-5
Switch Operating Rules for RADIUS	6-6
General RADIUS Setup Procedure	6-7

Configuring the Switch for RADIUS Authentication	6-8
Outline of the Steps for Configuring RADIUS Authentication	6-9
1. Configure Authentication for the Access Methods You Want RADIUS To Protect	6-10
2. Enable the (Optional) Access Privilege Option	6-12
3. Configure the Switch To Access a RADIUS Server	6-13
4. Configure the Switch's Global RADIUS Parameters	6-15
Using SNMP To View and Configure Switch Authentication Features 6-19	
Changing and Viewing the SNMP Access Configuration	6-20
Local Authentication Process	6-22
Controlling Web Browser Interface Access	6-23
Configuring RADIUS Authorization	6-24
Overview	6-24
Commands Authorization Type	6-24
Enabling Authorization with the CLI	6-25
Showing Authorization Information	6-26
Configuring the RADIUS Server	6-26
Using Vendor Specific Attributes (VSAs)	6-26
Example Configuration on Cisco Secure ACS for MS Windows	6-28
Example Configuration Using FreeRADIUS	6-30
Configuring RADIUS Accounting	6-32
Operating Rules for RADIUS Accounting	6-33
Steps for Configuring RADIUS Accounting	6-34
1. Configure the Switch To Access a RADIUS Server	6-35
2. Configure Accounting Types and the Controls for Sending Reports to the RADIUS Server	6-36
3. (Optional) Configure Session Blocking and Interim Updating Options	6-38
Viewing RADIUS Statistics	6-40
General RADIUS Statistics	6-40
RADIUS Authentication Statistics	6-42
RADIUS Accounting Statistics	6-43
Changing RADIUS-Server Access Order	6-44
Messages Related to RADIUS Operation	6-47

7 Configuring RADIUS Server Support for Switch Services

Contents	7-1
Overview	7-2
Configuring the RADIUS Server for Per-Port CoS and Rate-Limiting Services	7-3
Viewing the Currently Active Per-Port CoS and Rate-Limiting Configuration Specified by a RADIUS Server	7-4
Configuring and Using RADIUS-Assigned Access Control Lists ...	7-8
Introduction	7-8
Terminology	7-8
Overview of RADIUS-Assigned, Dynamic Port ACLs	7-11
Contrasting Dynamic and Static ACLs	7-13
How a RADIUS Server Applies a Dynamic Port ACL to a Switch Port ..	7-15
General ACL Features, Planning, and Configuration	7-16
The Packet-filtering Process	7-16
Operating Rules for Dynamic Port ACLs	7-17
Configuring an ACL in a RADIUS Server	7-18
Configuring ACE Syntax in RADIUS Servers	7-21
Configuration Notes	7-22
Configuring the Switch To Support Dynamic Port ACLs	7-24
Displaying the Current Dynamic Port ACL Activity on the Switch	7-25
Event Log Messages	7-28
Causes of Client Deauthentication Immediately After Authenticating	7-29
Monitoring Shared Resources	7-29

8 Configuring Secure Shell (SSH)

Contents	8-1
Overview	8-2
Terminology	8-3
Prerequisite for Using SSH	8-5

Public Key Formats	8-5
Steps for Configuring and Using SSH for Switch and Client Authentication	8-6
General Operating Rules and Notes	8-8
Configuring the Switch for SSH Operation	8-9
1. Assigning a Local Login (Operator) and Enable (Manager) Password	8-9
2. Generating the Switch's Public and Private Key Pair	8-10
3. Providing the Switch's Public Key to Clients	8-12
4. Enabling SSH on the Switch and Anticipating SSH Client Contact Behavior	8-15
5. Configuring the Switch for SSH Authentication	8-18
6. Use an SSH Client To Access the Switch	8-21
Further Information on SSH Client Public-Key Authentication .	8-22
Messages Related to SSH Operation	8-27

9 Configuring Secure Socket Layer (SSL)

Contents	9-1
Overview	9-2
Terminology	9-3
Prerequisite for Using SSL	9-5
Steps for Configuring and Using SSL for Switch and Client Authentication	9-5
General Operating Rules and Notes	9-6
Configuring the Switch for SSL Operation	9-7
1. Assigning a Local Login (Operator) and Enable (Manager)Password	9-7
2. Generating the Switch's Server Host Certificate	9-9
To Generate or Erase the Switch's Server Certificate with the CLI	9-10
Comments on certificate fields.	9-11
Generate a Self-Signed Host Certificate with the Web browser interface	9-13

Generate a CA-Signed server host certificate with the Web browser interface	9-15
3. Enabling SSL on the Switch and Anticipating SSL	
Browser Contact Behavior	9-17
Using the CLI interface to enable SSL	9-19
Using the web browser interface to enable SSL	9-19
Common Errors in SSL setup	9-21

10 Access Control Lists (ACLs)

Contents	10-1
Introduction	10-4
Overview of Options for Applying ACLs on the Switch	10-5
Static ACLS	10-5
Dynamic Port ACLs	10-5
Terminology	10-10
Overview	10-15
Types of IP ACLs	10-15
ACL Applications	10-15
RACL Applications	10-16
VACL Applications	10-18
Static Port ACL and Dynamic Port ACL Applications	10-19
Multiple ACLs on an Interface	10-20
Features Common to All ACL Applications	10-22
General Steps for Planning and Configuring ACLs	10-24
ACL Operation	10-26
Introduction	10-26
The Packet-filtering Process	10-27
Planning an ACL Application	10-30
IP Traffic Management and Improved Network Performance	10-30
Security	10-32
Guidelines for Planning the Structure of an ACL	10-32
ACL Configuration and Operating Rules	10-33
How an ACE Uses a Mask To Screen Packets for Matches	10-36

What Is the Difference Between Network (or Subnet) Masks and the Masks Used with ACLs?	10-36
Rules for Defining a Match Between a Packet and an Access Control Entry (ACE)	10-37
Configuring and Assigning an ACL	10-41
Overview	10-41
General Steps for Implementing ACLs	10-41
Options for Permit/Deny Policies	10-42
ACL Configuration Structure	10-42
Standard ACL Structure	10-43
Extended ACL Configuration Structure	10-45
ACL Configuration Factors	10-46
The Sequence of Entries in an ACL Is Significant	10-46
Allowing for the Implied Deny Function	10-48
A Configured ACL Has No Effect Until You Apply It to an Interface	10-48
You Can Assign an ACL Name or Number to an Interface Even if the ACL Does Not Exist in the Switch's Configuration	10-48
Using the CLI To Create an ACL	10-49
General ACE Rules	10-49
Using CIDR Notation To Enter the ACL Mask	10-50
Configuring Standard ACLs	10-51
Configuring Named, Standard ACLs	10-53
Creating Numbered, Standard ACLs	10-56
Configuring Extended ACLs	10-60
Configuring Named, Extended ACLs	10-62
Configuring Numbered, Extended ACLs	10-74
Adding or Removing an ACL Assignment On an Interface	10-81
Filtering Routed IP Traffic	10-81
Filtering IP Traffic Inbound on a VLAN	10-82
Filtering Inbound IP Traffic Per Port	10-84
Deleting an ACL	10-85
Editing an Existing ACL	10-86
Using the CLI To Edit ACLs	10-86
General Editing Rules	10-86

Sequence Numbering in ACLs	10-87
Inserting an ACE in an Existing ACL	10-88
Deleting an ACE from an Existing ACL	10-90
Resequencing the ACEs in an ACL	10-91
Attaching a Remark to an ACE	10-92
Operating Notes for Remarks	10-95
Displaying ACL Configuration Data	10-96
Display an ACL Summary	10-97
Display the Content of All ACLs on the Switch	10-98
Display the RAACL and VAACL Assignments for a VLAN	10-99
Display Static Port ACL Assignments	10-100
Displaying the Content of a Specific ACL	10-101
Display All ACLs and Their Assignments in the Routing Switch Startup-Config File and Running-Config File	10-103
Creating or Editing ACLs Offline	10-104
Creating or Editing an ACL Offline	10-104
The Offline Process	10-104
Example of Using the Offline Process	10-105
Enable ACL “Deny” Logging	10-109
Requirements for Using ACL Logging	10-109
ACL Logging Operation	10-110
Enabling ACL Logging on the Switch	10-111
General ACL Operating Notes	10-113

11 Configuring Advanced Threat Protection

Contents	11-1
Introduction	11-2
DHCP Snooping	11-3
Overview	11-3
Enabling DHCP Snooping	11-4
Enabling DHCP Snooping on VLANS	11-6
Configuring DHCP Snooping Trusted Ports	11-7
Configuring Authorized Server Addresses	11-8
Using DHCP Snooping with Option 82	11-8

Changing the Remote-id from a MAC to an IP Address	11-10
Disabling the MAC Address Check	11-10
The DHCP Binding Database	11-11
Operational Notes	11-12
Log Messages	11-13
Dynamic ARP Protection	11-15
Introduction	11-15
Enabling Dynamic ARP Protection	11-17
Configuring Trusted Ports	11-17
Adding an IP-to-MAC Binding to the DHCP Database	11-18
Configuring Additional Validation Checks on ARP Packets	11-19
Verifying the Configuration of Dynamic ARP Protection	11-20
Displaying ARP Packet Statistics	11-21
Monitoring Dynamic ARP Protection	11-21
Using the Instrumentation Monitor	11-22
Operating Notes	11-23
Configuring Instrumentation Monitor	11-24
Examples	11-25
Viewing the Current Instrumentation Monitor Configuration	11-26

12 Traffic/Security Filters and Monitors

Contents	12-1
Overview	12-2
Introduction	12-2
Filter Limits	12-3
Using Port Trunks with Filters	12-3
Filter Types and Operation	12-3
Source-Port Filters	12-4
Operating Rules for Source-Port Filters	12-4
Example	12-5
Named Source-Port Filters	12-6
Operating Rules for Named Source-Port Filters	12-6
Defining and Configuring Named Source-Port Filters	12-7
Viewing a Named Source-Port Filter	12-8

Using Named Source-Port Filters	12-9
Static Multicast Filters	12-15
Protocol Filters	12-16
Configuring Traffic/Security Filters	12-17
Configuring a Source-Port Traffic Filter	12-18
Example of Creating a Source-Port Filter	12-19
Configuring a Filter on a Port Trunk	12-19
Editing a Source-Port Filter	12-20
Configuring a Multicast or Protocol Traffic Filter	12-21
Filter Indexing	12-22
Displaying Traffic/Security Filters	12-23

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

Contents	13-1
Overview	13-3
Why Use Port-Based or User-Based Access Control?	13-3
General Features	13-3
User Authentication Methods	13-4
802.1X User-Based Access Control	13-4
802.1X Port-Based Access Control	13-5
Alternative To Using a RADIUS Server	13-6
Accounting	13-6
Terminology	13-6
General 802.1X Authenticator Operation	13-9
Example of the Authentication Process	13-9
VLAN Membership Priority	13-10
General Operating Rules and Notes	13-12
General Setup Procedure for 802.1X Access Control	13-14
Do These Steps Before You Configure 802.1X Operation	13-14
Overview: Configuring 802.1X Authentication on the Switch	13-15
Configuring Switch Ports as 802.1X Authenticators	13-16
1. Enable 802.1X Authentication on Selected Ports	13-17

A. Enable the Selected Ports as Authenticators and Enable the (Default) Port-Based Authentication	13-17
B. Specify User-Based Authentication or Return to Port-Based Authentication	13-18
Example: Configuring User-Based 802.1X Authentication	13-19
Example: Configuring Port-Based 802.1X Authentication	13-19
2. Reconfigure Settings for Port-Access	13-19
3. Configure the 802.1X Authentication Method	13-21
4. Enter the RADIUS Host IP Address(es)	13-22
5. Enable 802.1X Authentication on the Switch	13-23
6. Optional: Reset Authenticator Operation	13-23
7. Optional: Configure 802.1X Controlled Directions	13-24
Wake-on-LAN Traffic	13-24
Operating Notes	13-25
Example: Configuring 802.1X Controlled Directions	13-25
802.1X Open VLAN Mode	13-26
Introduction	13-26
VLAN Membership Priorities	13-27
Use Models for 802.1X Open VLAN Modes	13-28
Operating Rules for Authorized-Client and Unauthorized-Client VLANs	13-33
Setting Up and Configuring 802.1X Open VLAN Mode	13-37
802.1X Open VLAN Operating Notes	13-41
Option For Authenticator Ports: Configure Port-Security To Allow Only 802.1X-Authenticated Devices	13-42
Port-Security	13-43
Configuring Switch Ports To Operate As Supplicants for 802.1X Connections to Other Switches	13-44
Example	13-44
Supplicant Port Configuration	13-46
Displaying 802.1X Configuration, Statistics, and Counters	13-48
Show Commands for Port-Access Authenticator	13-48
Viewing 802.1X Open VLAN Mode Status	13-51
Show Commands for Port-Access Supplicant	13-55
How RADIUS/802.1X Authentication Affects VLAN Operation .	13-56

Operating Notes	13-60
Messages Related to 802.1X Operation	13-61

14 Configuring and Monitoring Port Security

Contents	14-1
Overview	14-3
Port Security	14-4
Basic Operation	14-4
Eavesdrop Protection	14-5
Blocking Unauthorized Traffic	14-5
Trunk Group Exclusion	14-6
Planning Port Security	14-7
Port Security Command Options and Operation	14-8
Port Security Display Options	14-8
Configuring Port Security	14-12
Retention of Static Addresses	14-18
MAC Lockdown	14-23
Differences Between MAC Lockdown and Port Security	14-25
MAC Lockdown Operating Notes	14-26
Deploying MAC Lockdown	14-27
MAC Lockout	14-31
Port Security and MAC Lockout	14-33
Web: Displaying and Configuring Port Security Features	14-34
Reading Intrusion Alerts and Resetting Alert Flags	14-34
Notice of Security Violations	14-34
How the Intrusion Log Operates	14-35
Keeping the Intrusion Log Current by Resetting Alert Flags	14-36
Menu: Checking for Intrusions, Listing Intrusion Alerts, and Resetting Alert Flags	14-37
CLI: Checking for Intrusions, Listing Intrusion Alerts, and Resetting Alert Flags	14-38
Using the Event Log To Find Intrusion Alerts	14-40
Web: Checking for Intrusions, Listing Intrusion Alerts, and Resetting Alert Flags	14-41

Operating Notes for Port Security	14-42
--	-------

15 Using Authorized IP Managers

Contents	15-1
Overview	15-2
Options	15-3
Access Levels	15-3
Defining Authorized Management Stations	15-4
Overview of IP Mask Operation	15-4
Menu: Viewing and Configuring IP Authorized Managers	15-5
CLI: Viewing and Configuring Authorized IP Managers	15-6
Listing the Switch's Current Authorized IP Manager(s)	15-6
Configuring IP Authorized Managers for the Switch	15-7
Web: Configuring IP Authorized Managers	15-9
Building IP Masks	15-9
Configuring One Station Per Authorized Manager IP Entry	15-9
Configuring Multiple Stations Per Authorized Manager IP Entry ..	15-10
Additional Examples for Authorizing Multiple Stations	15-12
Operating Notes	15-12

16 Key Management System

Contents	16-1
Overview	16-2
Terminology	16-2
Configuring Key Chain Management	16-3
Creating and Deleting Key Chain Entries	16-3
Assigning a Time-Independent Key to a Chain	16-4
Assigning Time-Dependent Keys to a Chain	16-5

Index

Product Documentation

About Your Switch Manual Set

Note

For the latest version of all ProCurve switch documentation, including Release Notes covering recently added features, please visit the ProCurve Networking Web site at www.procurve.com, click on **Technical support**, and then click on **Product manuals (all)**.

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 VLANs, MSTP, QoS, and Meshing.
- *Multicast and Routing Guide*—Explains how to configure IGMP, PIM, IP routing, and VRRP features.
- *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 main product guide.

Software Feature Index

For the software manual set supporting your 3500yl/5400zl/6200yl switch model, this feature index indicates which manual to consult for information on a given software feature.

Premium Edge Software Features. For the ProCurve 3500yl and 5400zl switches, Premium Edge features can be acquired by purchasing the optional Premium Edge license and installing it on the Intelligent Edge version of these switches. (These features are automatically included on the ProCurve 6200yl switches.)

Intelligent Edge Software Features. These features are automatically included on the ProCurve 3500yl and 5400zl Intelligent Edge switches and on the 6200yl Premium Edge switch.

Premium Edge Software Features	Manual			
	Management and Configuration	Advanced Traffic Management	Multicast and Routing	Access Security Guide
OSPF			X	
PIM-DM (Dense Mode)			X	
PIM-SM (Sparse Mode)			X	
VRRP			X	

Intelligent Edge Software Features	Manual			
	Management and Configuration	Advanced Traffic Management	Multicast and Routing	Access Security Guide
802.1Q VLAN Tagging		X		
802.1X Port-Based Priority	X			
802.1X Multiple Authenticated Clients Per Port				X
ACLs				X

Intelligent Edge Software Features	Manual			
	Management and Configuration	Advanced Traffic Management	Multicast and Routing	Access Security Guide
AAA Authentication				X
Authorized IP Managers				X
Authorized Manager List (Web, Telnet, TFTP)				X
Auto MDIX Configuration	X			
BOOTP	X			
Config File	X			
Console Access	X			
Copy Command	X			
CoS (Class of Service)		X		
Debug	X			
DHCP Configuration	X			
DHCP Option 82			X	
DHCP Snooping				X
DHCP/Bootp Operation	X			
Diagnostic Tools	X			
Downloading Software	X			
Dynamic ARP Protection				X
Eavesdrop Protection				X
Event Log	X			
Factory Default Settings	X			
Flow Control (802.3x)	X			
File Management	X			
File Transfers	X			
Friendly Port Names	X			
Guaranteed Minimum Bandwidth (GMB)	X			

Intelligent Edge Software Features	Manual			
	Management and Configuration	Advanced Traffic Management	Multicast and Routing	Access Security Guide
GVRP		X		
Identity-Driven Management (IDM)		X		
IGMP			X	
Interface Access (Telnet, Console/Serial, Web)	X			
IP Addressing	X			
IP Routing			X	
Jumbo Packets	X			
LACP	X			
Link	X			
LLDP	X			
LLDP-MED	X			
MAC Address Management	X			
MAC Lockdown				X
MAC Lockout				X
MAC-based Authentication				X
Management VLAN		X		
Meshing		X		
Monitoring and Analysis	X			
Multicast Filtering				X
Multiple Configuration Files	X			
Network Management Applications (SNMP)	X			
OpenView Device Management	X			
Passwords and Password Clear Protection				X
ProCurve Manager (PCM)	X			
Ping	X			

Intelligent Edge Software Features	Manual			
	Management and Configuration	Advanced Traffic Management	Multicast and Routing	Access Security Guide
Port Configuration	X			
Port Monitoring		X		
Port Security				X
Port Status	X			
Port Trunking (LACP)	X			
Port-Based Access Control (802.1X)				X
Power over Ethernet (PoE)	X			
Protocol Filters				X
Protocol VLANs		X		
Quality of Service (QoS)		X		
RADIUS Authentication and Accounting				X
RADIUS-Based Configuration				X
Rate-Limiting	X			
RIP			X	
RMON 1,2,3,9	X			
Routing			X	
Routing - IP Static			X	
Secure Copy	X			
sFlow	X			
SFTP	X			
SNMPv3	X			
Software Downloads (SCP/SFTP, TFPT, Xmodem)	X			
Source-Port Filters				X
Spanning Tree (STP, RSTP, MSTP)		X		
SSHv2 (Secure Shell) Encryption				X

Intelligent Edge Software Features	Manual			
	Management and Configuration	Advanced Traffic Management	Multicast and Routing	Access Security Guide
SSL (Secure Socket Layer)				X
Stack Management (3500yl/6200yl switches only)		X		
Syslog	X			
System Information	X			
TACACS+ Authentication				X
Telnet Access	X			
TFTP	X			
Time Protocols (TimeP, SNTP)	X			
Traffic Mirroring	X			
Traffic/Security Filters				X
Troubleshooting	X			
Uni-Directional Link Detection (UDLD)	X			
UDP Forwarder			X	
USB Device Support	X			
Virus Throttling (Connection-Rate Filtering)				X
VLANs		X		
VLAN Mirroring (1 static VLAN)		X		
Voice VLAN		X		
Web Authentication RADIUS Support				X
Web-based Authentication				X
Web UI	X			
Xmodem	X			

Security Overview

Contents

Introduction	1-2
About This Guide	1-2
For More Information	1-2
Switch Access Security	1-3
Default Configuration Settings and Access Security	1-3
Secure File Transfers	1-6
Other Provisions for Management Access Security	1-7
Network Security Features	1-8
Access Control Lists (ACLs)	1-8
802.1X Access Control	1-8
Web and MAC Authentication	1-9
Secure Shell (SSH)	1-9
Secure Socket Layer (SSLv3/TLSv1)	1-10
Traffic/Security Filters	1-10
Port Security, MAC Lockdown, and MAC Lockout	1-10
Key Management System (KMS)	1-11
Advanced Threat Detection	1-12
BPDU Filtering and BPDU Protection	1-12
Connection-Rate Filtering Based On Virus-Throttling Technology	1-12
DHCP Snooping, Dynamic ARP Protection, and Instrumentation Monitor	1-12
Identity-Driven Manager (IDM)	1-13

Introduction

Before you connect your switch to a network, ProCurve strongly recommends that you review the Security Overview beginning on page 1-3. It outlines the potential threats for unauthorized switch and network access, and provides guidelines on how to use the various security features available on the switch to prevent such access. For more information on individual features, see the references provided.

About This Guide

This *Access Security Guide* describes how to configure security features on the following switch models:

- ProCurve Switch 5406zl
- ProCurve Switch 5412zl
- ProCurve Switch 3500yl-24G-PWR Intelligent Edge
- ProCurve Switch 3500yl-48G-PWR Intelligent Edge
- ProCurve Switch 6200yl-24G-mGBIC Premium Edge

Note

For an introduction to the standard conventions used in this guide, refer to the *Getting Started* chapter in the *Management and Configuration Guide* for your switch.

For More Information

For information on which product manual to consult for a specific software feature, refer to the “Software Feature Index” on page xx of this guide.

For the latest version of all ProCurve switch documentation, including Release Notes covering recently added features and other software topics, visit the ProCurve Networking web site at www.procurve.com, click on **Technical support**, and then click on **Product Manuals (all)**.

Switch Access Security

This section outlines provisions for protecting access to the switch's status information and configuration settings. ProCurve switches are designed as "plug and play" devices, allowing quick and easy installation in your network. However, when preparing the switch for network operation, ProCurve strongly recommends that you enforce a security policy to help ensure that the ease in getting started is not used by unauthorized persons as an opportunity for access and possible malicious actions. Since security incidents can originate with sources inside as well as outside of an organization, your access security provisions must protect against internal and external threats while preserving the necessary network access for authorized clients and users.

Default Configuration Settings and Access Security

In its default configuration, the switch is open to unauthorized access of various types. In addition to applying local passwords, ProCurve recommends that you consider using the switch's other security features to provide a more complete security fabric.

Switch management access is available through the following methods:

- Inbound Telnet access and Web-browser access
- SNMP access
- Front-Panel access (serial port access to the console, plus resets and clearing the password(s) or current configuration)

It is important to evaluate the level of management access vulnerability existing in your network and take steps to ensure that all reasonable security precautions are in place. This includes both configurable security options and physical access to the switch hardware.

Local Manager Password

In the default configuration, there is no password protection. Configuring a local Manager password is a fundamental step in reducing the possibility of unauthorized access through the switch's Web browser and console (CLI and Menu) interfaces. The Manager password can easily be set using the CLI **password manager** command, the Menu interface **Console Passwords** option, or the password options under the **Security** tab in the Web browser interface.

Inbound Telnet Access and Web Browser Access

The default remote management protocols enabled on the switch are plain text protocols, which transfer passwords in open or plain text that is easily captured. To reduce the chances of unauthorized users capturing your passwords, secure and encrypted protocols such as SSH and SSL must be used for remote access. This enables you to employ increased access security while still retaining remote client access.

- SSHv2 provides Telnet-like connections through encrypted and authenticated transactions.
- SSLv3/TLSv1 provides remote Web browser access to the switch via encrypted paths between the switch and management station clients capable of SSL/TLS operation.

(For information on SSH, refer to Chapter 8 “Configuring Secure Shell (SSH)”; for details on SSL, refer to Chapter 9, “Configuring Secure Socket Layer (SSL)”.)

Also, access security on the switch is incomplete without disabling Telnet and the standard Web browser access. Among the methods for blocking unauthorized access attempts using Telnet or the Web browser are the following two CLI commands:

- **no telnet-server:** This command blocks inbound Telnet access.
- **no web-management:** This command prevents use of the Web browser interface through http (port 80) server access.

If you choose not to disable Telnet and Web browser access, you may want to consider using RADIUS accounting to maintain a record of password-protected access to the switch. Refer to Chapter 6, “RADIUS Authentication and Accounting” in this guide.

SNMP Access (Simple Network Management Protocol)

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

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

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

SNMPv3 security options include:

- configuring device communities as a means for excluding management access by unauthorized stations
- configuring for access authentication and privacy
- reporting events to the switch CLI and to SNMP trap receivers
- restricting non-SNMPv3 agents to either read-only access or no access
- co-existing with SNMPv1 and v2c if necessary

SNMP Access to the Authentication Configuration MIB. Beginning with software release K.12.*xxx*, a management station running an SNMP networked device management application, such as ProCurve Manager Plus (PCM+) or HP OpenView, can access the switch's management information base (MIB) for read access to the switch's status and read/write access to the switch's authentication configuration (hpSwitchAuth). This means that the switch's default configuration now allows SNMP access to security settings in hpSwitchAuth.

**Note on SNMP
Access to
Authentication
MIB**

Downloading and booting from the K.12.*xxx* or greater software version for the first time enables SNMP access to the authentication configuration MIB (the default action). If SNMPv3 and other security safeguards are not in place, the switch's authentication configuration MIB is exposed to unprotected SNMP access and you should use the command [shown](#) below to disable this access.

If SNMP access to the hpSwitchAuth MIB is considered a security risk in your network, then you should implement the following security precautions when downloading and booting from software release K.12.*xxx* or greater:

- [If SNMP access to the authentication configuration \(hpSwitchAuth\) MIB described above is not desirable for your network, then immediately after downloading and booting from the K.12.*xxx* or greater software for the first time, use the following command to disable this feature:](#)

snmp-server mib hpswitchauthmib excluded

- If you choose to leave the authentication configuration MIB accessible, then you should do the following to help ensure that unauthorized workstations cannot use SNMP tools to access the MIB:
 - a. Configure SNMP version 3 management and access security on the switch.
 - b. Disable SNMP version 2c on the switch.

For details on this feature, refer to the section titled “Using SNMP To View and Configure Switch Authentication Features” on page 6-19.

For information on SNMP, refer to “Using SNMP Tools To Manage the Switch” in the chapter titled “Configuring for Network Management Applications” in the *Management and Configuration Guide* for your switch.

Front-Panel Access and Physical Security

Physical access to the switch allows the following:

- use of the console serial port (CLI and Menu interface) for viewing and changing the current configuration and for reading status, statistics, and log messages.
- use of the switch’s Clear and Reset buttons for these actions:
 - clearing (removing) local password protection
 - rebooting the switch
 - restoring the switch to the factory default configuration (and erasing any non-default configuration settings)

Keeping the switch in a locked wiring closet or other secure space helps to prevent unauthorized physical access. As additional precautions, you can do the following:

- Disable or re-enable the password-clearing function of the Clear button.
- Configure the Clear button to reboot the switch after clearing any local usernames and passwords.
- Modify the operation of the Reset+Clear button combination so that the switch reboots, but does not restore the switch’s factory default settings.
- Disable or re-enable password recovery.

For the commands used to implement the above actions, refer to the section titled “Front-Panel Security” on page 2-8.

Secure File Transfers

Secure Copy and SFTP provide a secure alternative to TFTP and auto-TFTP for transferring sensitive information such as configuration files and log information between the switch and other devices. For more on these features, refer to the section on “Using Secure Copy and SFTP” in the “File Transfers” appendix of the *Management and Configuration Guide* for your switch.

Other Provisions for Management Access Security

The following features can help to prevent unauthorized management access to the switch.

Authorized IP Managers

This feature uses IP addresses and masks to determine whether to allow management access to the switch across the network through the following :

- Telnet and other terminal emulation applications
- The switch's Web browser interface
- SNMP (with a correct community name)

For more information, refer to Chapter 15, "Using Authorized IP Managers".

Secure Management VLAN

This feature creates an isolated network for managing the ProCurve switches that offer this feature. When a secure management VLAN is enabled, CLI, Menu interface, and Web browser interface access is restricted to ports configured as members of the VLAN. For more information, refer to the chapter titled "Static Virtual LANs (VLANs)" in the *Advanced Traffic Management Guide*.

TACACS+ Authentication

This application uses a central server to allow or deny access to TACACS-aware devices in your network. TACACS+ uses username/password sets with associated privilege levels to grant or deny access through either the switch's serial (console) port or remotely, with Telnet. If the switch fails to connect to a TACACS+ server for the necessary authentication service, it defaults to its own locally configured passwords for authentication control. TACACS+ allows both login (read-only) and enable (read/write) privilege level access. For more information, refer to Chapter 5, "TACACS+ Authentication".

RADIUS Authentication

For each authorized client, RADIUS can be used to authenticate operator or manager access privileges on the switch via the serial port (CLI and Menu interface), Telnet, SSH, and Secure FTP/Secure Copy (SFTP/SCP) access methods. Refer to Chapter 6, "RADIUS Authentication and Accounting".

ACLs for Management Access Protection

ACLs can also be configured to protect management access by blocking inbound IP traffic that has the switch itself as the destination IP address. (Refer to "Access Control Lists (ACLs)" in the next section.)

Network Security Features

This section outlines features for protecting access through the switch to the network. For more detailed information, see the indicated chapters.

Access Control Lists (ACLs)

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:** Eliminating unwanted IP, TCP, or UDP traffic by filtering packets where they enter or leave the switch on specific interfaces.

ACLs can filter traffic to or from a host, a group of hosts, or entire subnets. For details on how to apply ACLs in a network populated with ProCurve switches that support ACLs, see Chapter 10, “Access Control Lists (ACLs)”.

Note on ACL Security Use

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.*

802.1X Access Control

This feature provides port-based or user-based authentication through a RADIUS server to protect the switch from unauthorized access and to enable the use of RADIUS-based user profiles to control client access to network services. Included in the general features are the following:

- user-based access control supporting up to 32 authenticated clients per port
- port-based access control allowing authentication by a single client to open the port
- switch operation as a supplicant for point-to-point connections to other 802.1X-compliant ProCurve switches

For more information, refer to Chapter 13 “Configuring Port-Based and User-Based Access Control (802.1X)”.

Web and MAC Authentication

These options are designed for application on the edge of a network to provide port-based security measures for protecting private networks and the switch itself from unauthorized access. Because neither method requires clients to run any special supplicant software, both are suitable for legacy systems and temporary access situations where introducing supplicant software is not an attractive option. Both methods rely on using a RADIUS server for authentication. This simplifies access security management by allowing you to control access from a master database in a single server. It also means the same credentials can be used for authentication, regardless of which switch or switch port is the current access point into the LAN. Web authentication uses a web page login to authenticate users for access to the network. MAC authentication grants access to a secure network by authenticating device MAC addresses for access to the network. For more information, refer to Chapter 4, “Web and MAC Authentication”.

Secure Shell (SSH)

SSH provides Telnet-like functions through encrypted, authenticated transactions of the following types:

- **client public-key authentication:** uses one or more public keys (from clients) that must be stored on the switch. Only a client with a private key that matches a stored public key can gain access to the switch.
- **switch SSH and user password authentication:** this option is a subset of the client public-key authentication, and is used if the switch has SSH enabled without a login access configured to authenticate the client’s key. In this case, the switch authenticates itself to clients, and users on SSH clients then authenticate themselves to the switch by providing passwords stored on a RADIUS or TACACS+ server, or locally on the switch.
- **secure copy (SC) and secure FTP (SFTP):** By opening a secure, encrypted SSH session, you can take advantage of SC and SFTP to provide a secure alternative to TFTP for transferring sensitive switch information.

For more information on SSH, refer to Chapter 8, “Configuring Secure Shell (SSH)”. For more on SC and SFTP, refer to the section titled “Using Secure Copy and SFTP” in the “File Transfers” appendix of the *Management and Configuration Guide* for your switch.

Secure Socket Layer (SSLv3/TLSv1)

This feature includes use of Transport Layer Security (TLSv1) to provide remote web access to the switch via authenticated transactions and encrypted paths between the switch and management station clients capable of SSL/TLS operation. The authenticated type includes server certificate authentication with user password authentication. For more information, refer to Chapter 9, “Configuring Secure Socket Layer (SSL)”.

Traffic/Security Filters

These statically configured filters enhance in-band security (and improve control over access to network resources) by forwarding or dropping inbound network traffic according to the configured criteria. Filter options include:

- **source-port filters:** Inbound traffic from a designated, physical source-port will be forwarded or dropped on a per-port (destination) basis.
- **multicast filters:** Inbound traffic having a specified multicast MAC address will be forwarded to outbound ports or dropped on a per-port (destination) basis.
- **protocol filters:** Inbound traffic having the selected frame (protocol) type will be forwarded or dropped on a per-port (destination) basis.

For details, refer to Chapter 12, “Traffic/Security Filters and Monitors”.

Port Security, MAC Lockdown, and MAC Lockout

The features listed below provide device-based access security in the following ways:

- **Port security:** Enables configuration of 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. Some switch models also include eavesdrop prevention in the port security feature.
- **MAC lockdown:** This “static addressing” feature is used as an alternative to port security to prevent station movement and MAC address “hijacking” by allowing a given MAC address to use only one assigned port on the switch. MAC lockdown also restricts the client device to a specific VLAN.
- **MAC lockout:** This feature enables blocking of a specific MAC address so that the switch drops all traffic to or from the specified address.

Precedence of Security Options. 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.)

For more information, refer to Chapter 14, “Configuring and Monitoring Port Security”.

Key Management System (KMS)

KMS is available in several ProCurve switch models and is designed to configure and maintain key chains for use with KMS-capable routing protocols that use time-dependent or time-independent keys. (A key chain is a set of keys with a timing mechanism for activating and deactivating individual keys.) KMS provides specific instances of routing protocols with one or more Send or Accept keys that must be active at the time of a request.

For more information, refer to Chapter 16, “Key Management System”.

Advanced Threat Detection

Advanced threat detection covers a range of features used to detect anomalous traffic on the switch and take mitigating action against network attacks.

BPDU Filtering and BPDU Protection

Protects the network from denial-of-service attacks that use spoofing BPDUs by dropping incoming BPDU frames and/or blocking traffic through a port. For more information, see “Configuring BPDU Filtering” and “Configuring BPDU Protection” in the chapter titled “Multiple Instance Spanning-Tree Operation” in the *Advanced Traffic Management Guide* for your switch.

Connection-Rate Filtering Based On Virus-Throttling Technology

While not specifically a tool for controlling network access, this feature does help to protect the network from attack and is recommended for use on the network edge. It is primarily focused on the class of worm-like malicious code that tries to replicate itself by taking advantage of weaknesses in network applications behind unsecured ports. In this case, the malicious code tries to create a large number of outbound IP connections on an interface in a short time. Connection-Rate filtering detects hosts that are generating IP traffic that exhibits this behavior, and causes the switch to generate warning messages and (optionally) to either throttle or drop all IP traffic from the offending hosts. Refer to Chapter 3, “Virus Throttling” for details.

DHCP Snooping, Dynamic ARP Protection, and Instrumentation Monitor

These features provide the following additional protections for your network:

- **DHCP Snooping:** Protects your network from common DHCP attacks, such as address spoofing and repeated address requests.
- **Dynamic ARP Protection:** Protects your network from ARP cache poisoning.
- **Instrumentation Monitor.** Helps identify a variety of other common attacks by generating alerts for detected anomalies on the switch.

Refer to Chapter 11, “Configuring Advanced Threat Protection” for details.

Identity-Driven Manager (IDM)

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

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

- approve or deny access at the edge of the network instead of in the core;
- distinguish among different users and what each is authorized to do;
- configure guest access without compromising internal security.

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

- authorized user identity
- authorized device identity (MAC address)
- software running on the device
- physical location in the network
- time of day

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

For more information on IDM, visit the ProCurve Web site at www.procurve.com, and click on **Products and Solutions**, then **Identity Driven Manager** (under **Network Management**).

— This page is intentionally unused —

Configuring Username and Password Security

Contents

Overview	2-2
Configuring Local Password Security	2-5
Menu: Setting Passwords	2-5
CLI: Setting Passwords and Usernames	2-7
Web: Setting Passwords and Usernames	2-8
SNMP: Setting Passwords and Usernames	2-8
Front-Panel Security	2-8
When Security Is Important	2-9
Front-Panel Button Functions	2-10
Clear Button	2-10
Reset Button	2-11
Restoring the Factory Default Configuration	2-11
Configuring Front-Panel Security	2-12
Disabling the Clear Password Function of the Clear Button on the Switch's Front Panel	2-14
Re-Enabling the Clear Button on the Switch's Front Panel and Setting or Changing the "Reset-On-Clear" Operation	2-16
Changing the Operation of the Reset+Clear Combination	2-17
Password Recovery	2-18
Disabling or Re-Enabling the Password Recovery Process	2-18
Password Recovery Process	2-20

Overview

Feature	Default	Menu	CLI	Web
Set Usernames	none	—	—	page 2-8
Set a Password	none	page 2-5	page 2-7	page 2-8
Delete Password Protection	n/a	page 2-6	page 2-7	page 2-8
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.

Notes

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.

Beginning with software release K.12.*xxx*, usernames and passwords for Manager and Operator access can also be configured using SNMP. For more information, refer to “Using SNMP To View and Configure Switch Authentication Features” on page 6-19.

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 an inactivity timer. This causes the console session to end after the specified period of inactivity, thus giving you added security against unauthorized console access. You can use either of the following to set the inactivity timer:

- **Menu Interface:** System Information screen (Select “2. Switch Configuration.”)
- **CLI:** Use the **console inactivity-timer < 0 | 1 | 5 | 10 | 15 | 20 | 30 | 60 | 120 >**

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.

If the switch has a password for both the Manager and Operator levels, and neither is entered correctly in response to the switch's password prompt, 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 chapter covers how to:

- Set passwords
- Delete passwords
- Recover from a lost password
- Maintain front-panel security

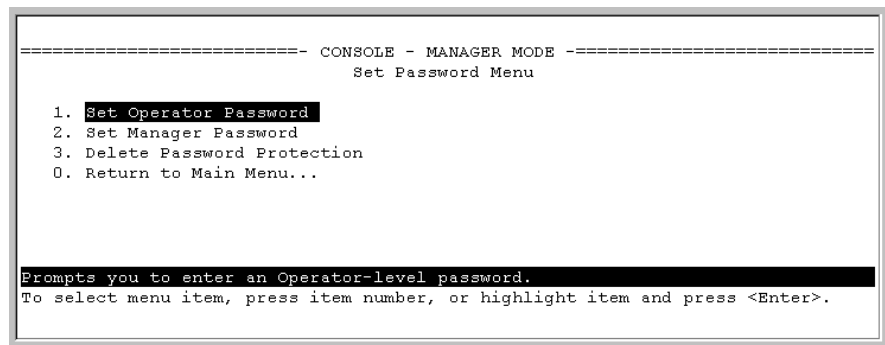
Configuring Local Password Security

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



```
----- CONSOLE - MANAGER MODE -----  
Set Password Menu  
  
1. Set Operator Password  
2. Set Manager Password  
3. Delete Password Protection  
0. Return to Main Menu...  
  
Prompts you to enter an Operator-level password.  
To select menu item, press item number, or highlight item and press <Enter>.
```

Figure 2-1. The Set Password Screen

2. To set a new password:
 - a. Select **Set Manager Password** or **Set Operator Password**. You will then be prompted with **Enter new password**.
 - b. Type a password of up to 16 ASCII characters with no spaces and press **[Enter]**. (Remember that passwords are case-sensitive.)
 - c. When prompted with **Enter new password again**, retype the new 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)#
```

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]**.

SNMP: Setting Passwords and Usernames

Beginning with software release K.12.*xxx*, usernames and passwords for Manager and Operator access can also be configured using SNMP. For more information, refer to “Using SNMP To View and Configure Switch Authentication Features” on page 6-19.

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.

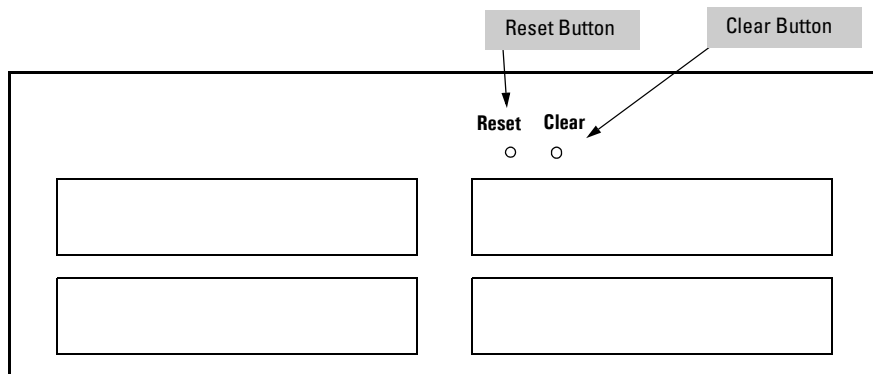


Figure 2-4. Front-Panel Button Locations on a ProCurve Series 5400zl Switch

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.



4. When the Test LED to the right of the Clear button begins flashing, release the Clear button.



It can take approximately 20-25 seconds for the switch to reboot. This process restores the switch configuration to the factory default settings.

Configuring Front-Panel Security

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

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

- Configure the Clear button to reboot the switch after clearing any local usernames and passwords. This provides an immediate, visual means (plus an Event Log message) for verifying that any usernames and passwords in the switch have been cleared.
- Modify the operation of the Reset+Clear combination (page 2-11) so that the switch still reboots, but does *not* restore the switch's factory default configuration settings. (Use of the Reset button alone, to simply reboot the switch, is not affected.)
- Disable or re-enable Password Recovery.

Syntax: show front-panel-security

Displays the current front-panel-security settings:

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

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

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

Password Recovery: Shows whether the switch is configured with the ability to recover a lost password. (Refer to “Password Recovery Process” on page 2-20.) (Default: **Enabled**.)

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

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

```
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 2-11.

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

```
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 2-11. Although the Clear button does not erase passwords when disabled, you can still use it with the Reset button (Reset+Clear) to restore the switch to its factory default configuration. You can then get access to the switch to set a new password.*

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


```

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 2-11 replaces the switch’s current startup-config file with the factory-default startup-config file, then reboots the switch, and removes local password protection. *This means that anyone who has physical access to the switch could use this button combination to replace the switch’s current configuration with the factory-default configuration, and render the switch accessible without the need to input a username or password.* You can use the **factory-reset** command to prevent the Reset+Clear combination from being used for this purpose.

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

Disables or re-enables the following functions associated with using the Reset+Clear buttons in the combination described under “Restoring the Factory Default Configuration” on page 2-11:

- *Replacing the current startup-config file with the factory-default startup-config file*
- *Clearing any local usernames and passwords configured on the switch*

(Default: Both functions enabled.)

Notes: *The Reset+Clear button combination always reboots the switch, regardless of whether the “no” form of the command has been used to disable the above two functions. Also, if you disable **factory-reset**, you cannot disable the **password-recovery** option, and the reverse.*

```
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 2-11) to restore the switch to its factory default configuration.

Note: To disable **password-recovery**:

- You must have physical access to the front panel of the switch.
- The **factory-reset** parameter must be enabled (the default).

(Default: Enabled.)

Steps for Disabling Password-Recovery.

1. Set the CLI to the global interface context.
2. Use **show front-panel-security** to determine whether the factory-reset parameter is enabled. If it is disabled, use the **front-panel-security factory-reset** command to enable it.
3. Press and release the Clear button on the front panel of the switch.
4. Within 60-seconds of pressing 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 2-11. This can disrupt network operation and make it necessary to temporarily disconnect the switch from the network to prevent unauthorized access and other problems while it is being reconfigured.

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

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

Virus Throttling

Contents

Overview of Connection-Rate Filtering	3-3
Features and Benefits	3-4
General Operation	3-5
Filtering Options	3-5
Sensitivity to Connection Rate Detection	3-5
Application Options	3-6
Operating Rules	3-7
Unblocking a Currently Blocked Host	3-7
General Configuration Guidelines	3-8
For a network that is relatively attack-free:	3-8
For a network that appears to be under significant attack:	3-9
Configuring Connection-Rate Filtering	3-10
Global and Per-Port Configuration	3-10
Enabling Connection-Rate Filtering and Configuring Sensitivity ...	3-11
Configuring the Per-Port Filtering Mode	3-12
Example of a Basic Connection-Rate Filtering Configuration ..	3-13
Viewing and Managing Connection-Rate Status	3-15
Viewing Connection-Rate Configuration	3-15
Listing Currently-Blocked Hosts	3-17
Unblocking Currently-Blocked Hosts	3-18
Configuring and Applying Connection-Rate ACLs	3-19
Connection-Rate ACL Operation	3-20
Configuring a Connection-Rate ACL Using	
Source IP Address Criteria	3-21
Configuring a Connection-Rate ACL Using UDP/TCP Criteria	3-23
Applying Connection-Rate ACLs	3-26
Using CIDR Notation To Enter the ACE Mask	3-26

Example of Using an ACL in a Connection-Rate Configuration	3-27
Connection-Rate ACL Operating Notes	3-29
Connection-Rate Log and Trap Messages	3-31

Overview of Connection-Rate Filtering

Feature	Default	Page Ref
Global Configuration and Sensitivity	Disabled	3-11
Per-Port Configuration	None	3-12
Listing and Unblocking Blocked Hosts	n/a	3-17
Viewing the Current Configuration	n/a	3-15
Configuring Connection-Rate ACLs	None	3-19

The spread of malicious agents in the form of worms exhibiting worm behavior has severe implications for network performance. Damage can be as minimal as slowing down a network with excessive, unwanted traffic, or as serious as putting attacker-defined code on a system to cause any type of malicious damage that an authorized user could do.

Current methods to stop the propagation of malicious agents rely on use of signature recognition to prevent hosts from being infected. However, the latency between the introduction of a new virus or worm into a network and the implementation and distribution of a signature-based patch can be significant. Within this period, a network can be crippled by the abnormally high rate of traffic generated by infected hosts.

Connection-rate filtering based on virus throttling technology is recommended for use on the edge of a network. It is primarily concerned with the class of worm-like malicious code that tries to replicate itself by using vulnerabilities on other hosts (that is, weaknesses in network applications behind unsecured ports). Agents of this variety operate by choosing a set of hosts to attack based on an address range (sequential or random) that is exhaustively searched, either by blindly attempting to make connections by rapidly sending datagrams to the address range, or by sending individual ICMP ping messages to the address range and listening for replies.

Connection-rate filtering exploits the network behavior of malicious code that tries to create a large number of outbound IP connections on a routed interface in a short time. When a host exhibits this behavior, warnings can be sent, and connection requests can be either throttled or dropped to minimize the barrage of subsequent traffic from the host. When enabled on the switch, connection-rate filtering can help reduce the impact of worm-like malicious code and give system administrators more time to isolate and eradicate the threat. Thus, while traditional worm and virus-signature updates will still need to be deployed to hosts, the network remains functional and the overall distribution of the malicious code is limited.

Features and Benefits

Connection-rate filtering is a countermeasure tool you can use in your incident-management program to help detect and manage worm-type IT security threats received in inbound IP traffic. Major benefits of this tool include:

- Behavior-based operation that does not require identifying details unique to the code exhibiting the worm-like operation.
- Handles unknown worms.
- Needs no signature updates.
- Protects network infrastructure by slowing or stopping IP traffic from hosts exhibiting high connection-rate behavior.
- Allows network and individual switches to continue to operate, even when under attack.
- Provides Event Log and SNMP trap warnings when worm-like behavior is detected
- Gives IT staff more time to react before the threat escalates to a crisis.

Note

When configured on a port, connection-rate filtering is triggered by IPv4 traffic received inbound with a relatively high rate of IP connection attempts.

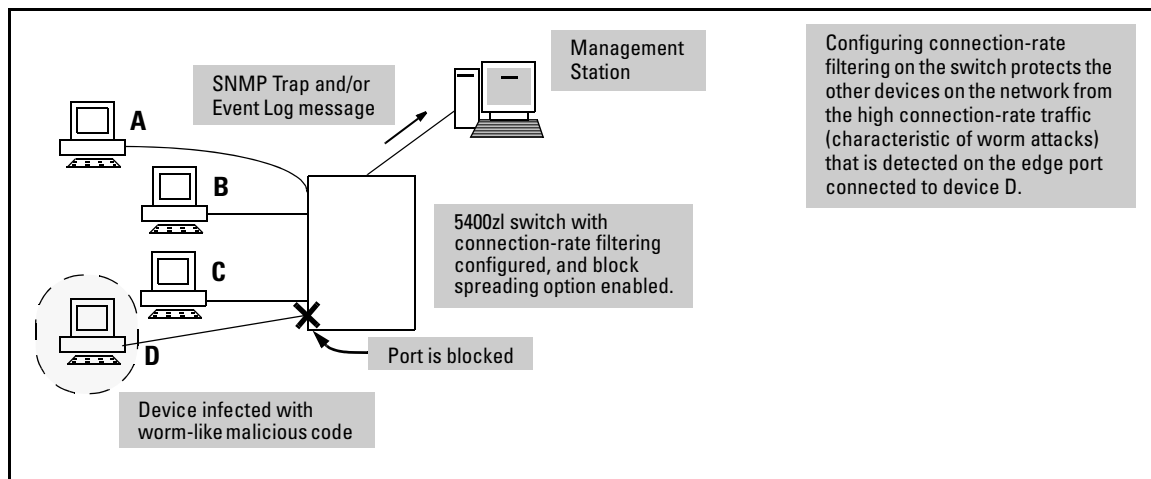


Figure 3-1. Example of Protecting a Network from Agents Using a High IP Connection Rate To Propagate

General Operation

Connection-rate filtering enables notification of worm-like behavior detected in inbound IP traffic and, depending on how you configure the feature, also throttles or blocks such traffic. This feature also provides a method for allowing legitimate, high connection-rate traffic from a given host while still protecting your network from possibly malicious traffic from other hosts.

Filtering Options

In the default configuration, connection-rate filtering is disabled. When enabled on a port, connection-rate filtering monitors inbound IP traffic for a high rate of connection requests from any given host on the port. If a host appears to exhibit the worm-like behavior of attempting to establish a large number of outbound IP connections in a short period of time, the switch responds in one of the following ways, depending on how connection-rate filtering is configured:

- **Notify only** (of potential attack): While the apparent attack continues, the switch generates an Event Log notice identifying the offending host's source IP address and (if a trap receiver is configured on the switch) a similar SNMP trap notice).
- **Throttle**: In this case, the switch temporarily blocks inbound IP traffic from the offending host source IP address for a "penalty" period and generates an Event Log notice of this action and (if a trap receiver is configured on the switch) a similar SNMP trap notice. When the "penalty" period expires the switch re-evaluates the traffic from the host and continues to block this traffic if the apparent attack continues. (During the re-evaluation period, IP traffic from the host is allowed.)
- **Block**: This option blocks all IP traffic from the host. When a block occurs, the switch generates an Event Log notice and (if a trap receiver is configured on the switch) a similar SNMP trap notice. Note that a network administrator must explicitly re-enable a host that has been previously blocked.

Sensitivity to Connection Rate Detection

The switch includes a global sensitivity setting that enables adjusting the ability of connection-rate filtering to detect relatively high instances of connection-rate attempts from a given source.

Application Options

For the most part, normal network traffic is distinct from the traffic exhibited by malicious agents. However, when a legitimate network host generates multiple connections in a short period of time, connection-rate filtering may generate a “false positive” and treat the host as an infected client. Lowering the sensitivity or changing the filter mode may reduce the number of false positives. Conversely, relaxing filtering and sensitivity provisions lowers the switch’s ability to detect worm-generated traffic in the early stages of an attack, and should be carefully investigated and planned to ensure that a risky vulnerability is not created. As an alternative, you can use connection-rate ACLs (*access control lists*) or selective enabling to allow legitimate traffic.

Selective Enable. This option involves applying connection-rate filtering only to ports posing a significant risk of attack. For ports that are reasonably secure from attack, then there may be little benefit in configuring them with connection-rate filtering.

Connection-Rate ACLs. The basic connection-rate filtering policy is configured per-port as **notify-only**, **throttle**, and **block**. A connection-rate ACL creates exceptions to these per-port policies by creating special rules for individual hosts, groups of hosts, or entire subnets. Thus, you can adjust a connection-rate filtering policy to create and apply an exception to configured filters on the ports in a VLAN. Note that connection-rate ACLs are useful only if you need to exclude inbound traffic from your connection-rate filtering policy. For example, a server responding to network demand may send a relatively high number of legitimate connection requests. This can generate a false positive by exhibiting the same elevated connection-rate behavior as a worm. Using a connection-rate ACL to apply an exception for this server allows you to exclude the trusted server from connection-rate filtering and thereby keep the server running without interruption.

Note

Use connection-rate ACLs only when you need to exclude an IP traffic source (including traffic with specific UDP or TCP criteria) from a connection-rate filtering policy. Otherwise, the ACL is not necessary.

Operating Rules

- Connection-rate filtering is triggered by inbound IP traffic exhibiting high rates of IP connections to new hosts. After connection-rate filtering has been triggered on a port, all traffic from the suspect host is subject to the configured connection-rate policy (**notify-only**, **throttle**, or **block**).
- When connection-rate filtering is configured on a port, the port cannot be added to, or removed from, a port trunk group. Before this can be done, connection-rate filtering must be disabled on the port.
- Where the switch is throttling or blocking inbound IP traffic from a host, any outbound traffic destined for that host is still permitted.
- Once a throttle has been triggered on a port—temporarily blocking inbound IP traffic—it cannot be undone during operation: the penalty period must expire before traffic will be allowed from the host.

Unblocking a Currently Blocked Host

A host blocked by connection-rate filtering remains blocked until explicitly unblocked by one of the following methods:

- Using the **connection-rate-filter unblock** command (page 3-17).
- Rebooting the switch.
- Disabling connection-rate filtering using the **no connection-rate-filter** command.
- Deleting a VLAN removes blocks on any hosts on that VLAN.

Note

Changing a port setting from **block** to **throttle**, **notify-only**, or to **no filter connection-rate**, does not unblock a currently blocked host. Similarly, applying a connection-rate ACL will not unblock a currently blocked host. Refer to the above list for the correct methods to use to unblock a host.

General Configuration Guidelines

As stated earlier, connection-rate filtering is triggered only by inbound IP traffic generating a relatively high number of new IP connection requests from the same host.

For a network that is relatively attack-free:

1. Enable **notify-only** mode on the ports you want to monitor.
2. Set global sensitivity to **low**.
3. If SNMP trap receivers are available in your network, use the **snmp-server** command to configure the switch to send SNMP traps.
4. Monitor the Event Log or (if configured) the available SNMP trap receivers to identify hosts exhibiting high connection rates.
5. Check any hosts that exhibit relatively high connection rate behavior to determine whether malicious code or legitimate use is the cause of the behavior.
6. Hosts demonstrating high, but legitimate connection rates, such as heavily used servers, may trigger a connection-rate filter. Configure connection rate ACLs to create policy exceptions for trusted hosts. (Exceptions can be configured for these criteria:
 - A single source host or group of source hosts
 - A source subnet
 - Either of the above with TCP or UDP criteria

(For more on connection rate ACLs, refer to “Application Options” on page 3-6.)

7. Increase the sensitivity to **Medium** and repeat steps 5 and 6.

Note

On networks that are relatively infection-free, sensitivity levels above **Medium** are not recommended.)

8. (Optional.) Enable **throttle** or **block** mode on the monitored ports.

Note

On a given VLAN, to unblock the hosts that have been blocked by the connection-rate feature, use the **vlan < vid > connection-rate filter unblock** command.

9. Maintain a practice of carefully monitoring the Event Log or configured trap receivers for any sign of high connectivity-rate activity that could indicate an attack by malicious code. (Refer to “Connection-Rate Log and Trap Messages” on page 3-31.)

For a network that appears to be under significant attack:

The steps are similar to the general steps for a network that is relatively attack free. The major difference is in policies suggested for managing hosts exhibiting high connection rates. This allows better network performance for unaffected hosts and helps to identify hosts that may require updates or patches to eliminate malicious code.

1. Configure connection-rate filtering to **throttle** on all ports.
2. Set global sensitivity to **medium**.
3. If SNMP trap receivers are available in your network, use the **snmp-server** command to configure the switch to send SNMP traps.
4. Monitor the Event Log or the available SNMP trap receivers (if configured on the switch) to identify hosts exhibiting high connection rates.
5. Check any hosts that exhibit relatively high connection rate behavior to determine whether malicious code or legitimate use is the cause of the behavior.
6. On hosts you identify as needing attention to remove malicious behavior:
 - To immediately halt an attack from a specific host, group of hosts, or a subnet, use the per-port block mode on the appropriate port(s).
 - After gaining control of the situation, you can use connection-rate ACLs to more selectively manage traffic to allow receipt of normal routed traffic from reliable hosts.

Configuring Connection-Rate Filtering

Command	Page
Global and Per-Port Configuration	
connection-rate-filter sensitivity < low medium high aggressive >	3-11
filter connection-rate < <i>port-list</i> > < notify-only throttle block >	3-12
show connection-rate-filter < blocked-host >	
Unblocking Hosts	
connection-rate-filter unblock	3-18

Note

As stated previously, connection-rate filtering is triggered by inbound IP traffic exhibiting a relatively high incidence of IP connection attempts from a single source.

Global and Per-Port Configuration

Use the commands in this section to enable connection-rate filtering on the switch and to apply the filtering on a per-port basis. (You can use the ACL commands in the next section to adjust a filter policy on a per-vlan basis to avoid filtering traffic from specific, trusted source addresses.)

Enabling Connection-Rate Filtering and Configuring Sensitivity

Syntax: connection-rate-filter sensitivity < low | medium | high | aggressive >
no connection-rate-filter

This command:

- *Enables connection-rate filtering.*
- *Sets the global sensitivity level at which the switch interprets a given host's attempts to connect to a series of different devices as a possible attack by a malicious agent residing in the host.*

Options for configuring sensitivity include:

low: *Sets the connection-rate sensitivity to the lowest possible sensitivity, which allows a mean of 54 destinations in less than 0.1 seconds, and a corresponding penalty time for Throttle mode (if configured) of less than 30 seconds.*

medium: *Sets the connection-rate sensitivity to allow a mean of 37 destinations in less than 1 second, and a corresponding penalty time for Throttle mode (if configured) between 30 and 60 seconds.*

high: *Sets the connection-rate sensitivity to allow a mean of 22 destinations in less than 1 second, and a corresponding penalty time for Throttle mode (if configured) between 60 and 90 seconds.*

aggressive: *Sets the connection-rate sensitivity to the highest possible level, which allows a mean of 15 routed destinations in less than 1 second, and a corresponding penalty time for Throttle mode (if configured) between 90 and 120 seconds.*

*The **no connection-rate-filter** command disables connection-rate filtering on the switch.*

Note

The sensitivity settings configured on the switch determines the Throttle mode penalty periods as shown in Table 3-1 on page 3-12.

Configuring the Per-Port Filtering Mode

Syntax: filter connection-rate < port-list > < notify-only | throttle | block >
no filter connection-rate < port-list >

*Configures the per-port policy for responding to detection of a relatively high number of inbound, routed IP connection attempts from a given source. The level at which the switch detects such traffic depends on the sensitivity setting configured by the **connection-rate-filter sensitivity** command (page 3-11). (Note: You can use connection-rate ACLs to create exceptions to the configured filtering policy. See “Configuring and Applying Connection-Rate ACLs” on page 3-19.) The **no** form of the command disables connection-rate filtering on the ports in # < port-list >.*

notify-only: *If the switch detects a relatively high number of routed IP connection attempts from a specific host, **notify-only** generates an Event Log message. Sends a similar message to any SNMP trap receivers configured on the switch.*

throttle: *If the switch detects a relatively high number of routed IP connection attempts from a specific host, this option generates the **notify-only** messaging and also blocks all routed traffic inbound from the offending host for a penalty period. After the penalty period, the switch allows routed traffic from the offending host to resume, and re-examines the traffic. If the suspect behavior continues, the switch again blocks the routed traffic from the offending host and repeats the cycle. For the penalty periods, refer to table 3-1, below.*

block: *If the switch detects a relatively high number of routed IP connection attempts from a specific host, this option generates the **notify-only** messaging and also blocks all routed and switched traffic inbound from the offending host.*

Table 3-1. Throttle Mode Penalty Periods

Throttle Mode (Sensitivity)	Frequency of IP Connection Requests from the Same Source	Mean Number of New Destination Hosts in the Frequency Period	Penalty Period
Low	< 0.1 second	54	< 30 seconds
Medium	< 1.0 second	37	30 - 60 seconds
High	< 1.0 second	22	60 - 90 seconds
Aggressive	< 1.0 second	15	90 - 120 seconds

Example of a Basic Connection-Rate Filtering Configuration

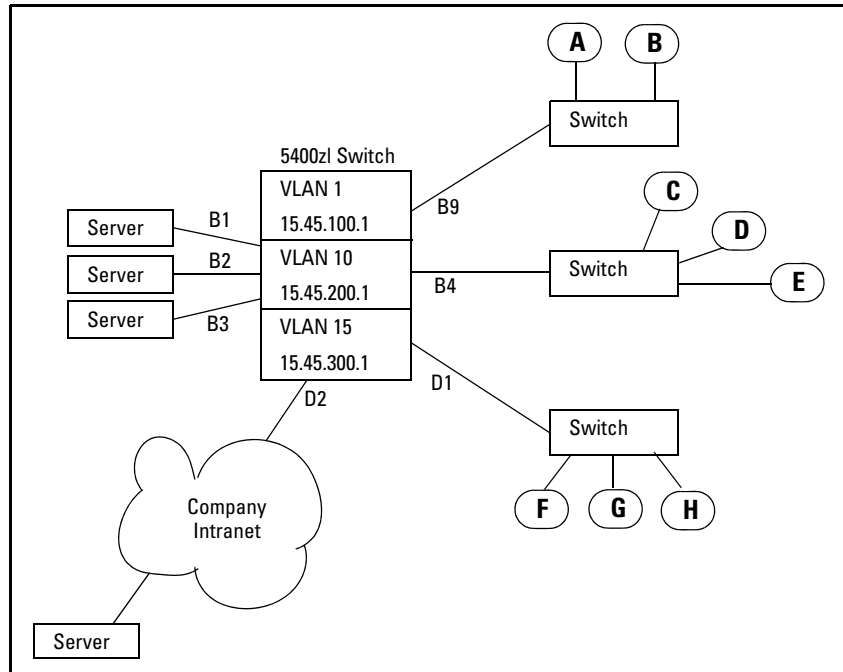


Figure 3-2. Sample Network

Basic Configuration. Suppose that in the sample network, the administrator wanted to enable connection-rate filtering and configure the following response to high connection-rate traffic on the switch:

- Ports B1 - B3: Throttle traffic from the transmitting host(s).
- Port B4: Respond with Notify-Only to identify the transmitting host(s).
- Ports B9, D1, and D2: Block traffic from the transmitting host(s).

Figure 3-3 illustrates the configuration steps and resulting startup-config file.

Virus Throttling

Configuring Connection-Rate Filtering

```
ProCurve(config)# connection-rate-filter sensitivity low
ProCurve(config)# filter connection-rate b1-b3 throttle
ProCurve(config)# filter connection-rate b4 notify-only
ProCurve(config)# filter connection-rate b9,d1-d2 block
ProCurve(config)# write mem
ProCurve(config)# show config
Startup configuration:
; J8697A Configuration Editor; Created on release #K.11.XX
hostname "ProCurve"
connection-rate-filter sensitivity low
module 2 type J8702A
module 4 type J8702A
ip routing
snmp-server community "public" Unrestricted
snmp-server host 15.45.200.75 "public"
vlan 1
  name "DEFAULT VLAN"
  untagged B5-B24
  ip address dhcp-bootp
  no untagged B1-B4,D1-D24
  ip proxy-arp
  exit
vlan 10
  name "VLAN10"
  untagged B1-B4
  no ip address
  ip proxy-arp
  exit
vlan 15
  name "VLAN15"
  untagged D1-D24
  no ip address
  ip proxy-arp
  exit
filter connection-rate B4 notify-only
filter connection-rate B1-B3 throttle
filter connection-rate B9,D1-D2 block
```

Enables connection-rate filtering and sets the sensitivity to "low".

Configures the desired responses to inbound, high connectivity-rate traffic on the various ports.

Indicates that connectivity-rate filtering is enabled at the "low" sensitivity setting.

Shows the per-port configuration for the currently enabled connectivity-rate filtering.

Figure 3-3. Example of a Basic Connection-Rate Configuration

Viewing and Managing Connection-Rate Status

The commands in this section describe how to:

- View the current connection-rate configuration
- List the currently blocked hosts
- Unblock currently blocked hosts

Viewing Connection-Rate Configuration

Use the following command to view the basic connection-rate configuration. If you need to view connection-rate ACLs and/or any other switch configuration details, use `show config` or `show running` (page 3-16).

Syntax: `show connection-rate-filter`

Displays the current global connection-rate status (enabled/disabled) and sensitivity setting, and the current per-port configuration. This command does not display the current (optional) connection-rate ACL configuration, if any.

```
ProCurve(config)# show connection-rate-filter
Connection Rate Filter Configuration
Global Status:      Enabled
Sensitivity:        Medium
Port                | Filter Mode
-----|-----
B13                 | NOTIFY-ONLY
B14                 | THROTTLE
B15                 | BLOCK
B16                 | BLOCK
```

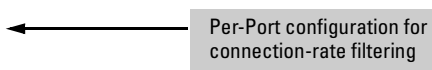


Figure 3-4. Example of Displaying the Connection-Rate Status, Sensitivity, and Per-Port Configuration

To view the complete connection-rate configuration, including any ACLs (page 3-19), use **show config** (for the startup-config file) or **show running** (for the running-config file). For example:

```
ProCurve (config)# show config
Startup configuration:
; J8697A Configuration Editor; Created on
hostname "ProCurve"
connection-rate-filter sensitivity medium
ip access-list connection-rate-filter "Sample"
  filter ip 13.28.234.180 0.0.15.255
  ignore ip 0.0.0.0 255.255.255.255
  exit
module 2 type J8161A
module 4 type J8161A
ip routing
logging 13.28.234.180
snmp-server community "public" Unrestricted
vlan 1
  name "DEFAULT_VLAN"
  untagged B1-B12,B19-B24,D1-D24
  no ip address
  no untagged B13-B18
  ip proxy-arp
  exit
vlan 15
  name "VLAN_15"
  untagged B13-B18
  ip address 13.28.234.181 255.255.240.0
  ip proxy-arp
  ip connection-rate-filter-access-group "Sample"
  exit
filter connection-rate B13 notify-only
filter connection-rate B14 throttle
filter connection-rate B15-B16 block
```

Entry showing that connection-rate-filtering is enabled and set to "medium" sensitivity.

Example of a connection-rate filtering ACL appearing in the configuration.

Example of a connection-rate filtering ACL appearing in a VLAN configuration.

Example of per-port connection-rate filtering policies appearing in the configuration.

Figure 3-5. Example of Connection-Rate Filtering Configuration in the Startup-Config File

Listing Currently-Blocked Hosts

Syntax: show connection-rate-filter < all-hosts | blocked-hosts | throttled-hosts >

all-hosts: Lists, by VLAN membership, all hosts currently detected in a throttling or blocking state, along with a state indicator.

throttled-hosts: Lists, by VLAN membership, the hosts currently in a throttling state due to connection-rate action.

blocked-hosts: Lists, by VLAN membership, the hosts currently blocked by connection-rate action.

```
ProCurve(config)# show connection-rate-filter all-hosts
```

VLAN ID	Source IP Address	Filter Mode
10	13.28.234.175	THROTTLE
10	13.28.234.179	THROTTLE
15	13.28.234.180	BLOCK

Figure 3-6. Example of Listing Hosts in Any Connection-Rate State

```
ProCurve(config)# show connection-rate-filter blocked-hosts
```

VLAN ID	Source IP Address
15	13.28.234.180

Figure 3-7. Example of Listing Hosts Blocked by Connection-Rate Filtering

Unblocking Currently-Blocked Hosts

If a host becomes blocked by triggering connection-rate filtering on a port configured to block high connection rates, the host remains blocked on all ports on the switch even if you change the per-port filtering configuration. (The source IP address block imposed by connection-rate filtering does not age-out.) This is to help prevent a malicious host from automatically regaining access to the network.

When a host becomes blocked the switch generates the following Event Log message and also sends a similar message to any configured SNMP trap receivers.

```
Src IP xxx.xxx.xxx.xxx blocked
```

Note

ProCurve recommends that, before you unblock a host that has been blocked by connection-rate filtering, you inspect the host with current antivirus tools and remove any malicious agents that pose a threat to your network.

If a trusted host frequently triggers connection-rate blocking with legitimate, high connection-rate traffic, then you may want to consider either changing the sensitivity level on the associated port or configuring a connection-rate ACL to create a filtering exception for the host.

Syntax: connection-rate-filter unblock < all | host | ip-addr >

all: *Unblocks all hosts currently blocked due to action by connection-rate filtering on ports where block mode has been configured.*

host < ip-addr >: *Unblocks the single host currently blocked due to action by connection-rate filtering on ports where block mode has been configured.*

ip-addr < mask > : *Unblocks traffic from any host in the specified subnet currently blocked due to action by connection-rate filtering on ports where block mode has been configured.*

*Note: There is also an option to unblock any host belonging to a specific VLAN using the **vlan <vid> connection-rate-filter unblock** command.*

Note

For a complete list of options for unblocking hosts, see page 3-7.

Configuring and Applying Connection-Rate ACLs

Command	Page
ip access-list connection-rate-filter < crf-list-name >	3-21, 3-23
< filter ignore > ip < any host < ip-addr > ip-addr < mask >>	3-21
< filter ignore > < udp tcp > < source > < options >	3-23
vlan < vid > ip access-group < crf-list-name > connection-rate-filter	

A host sending legitimate, routed traffic can trigger connection-rate filtering in some circumstances. If you can verify that such a host is indeed sending valid traffic and is not a threat to your network, you may want to configure a connection-rate ACL (access control list) that allows this traffic to bypass the configured connection-rate filtering.

A connection-rate Access Control List (ACL) is an optional tool that consists of one or more explicitly configured Access Control Entries (ACEs) used to specify whether to enforce the configured connection-rate policy on traffic from a particular source.

Use of connection-rate ACLs provides the option to apply exceptions to the configured connection-rate filtering policy. This enables you to allow legitimate traffic from a trusted source, and apply connection-rate filtering only to inbound traffic from untrusted sources. For example, where a connection-rate policy has been configured, you can apply a connection-rate ACL that causes the switch to bypass connection-rate policy filtering on traffic from:

- A trusted server exhibiting a relatively high IP connection rate due to heavy demand
- A trusted traffic source on the same port as other, untrusted traffic sources.

The criteria for an exception can include the source IP address of traffic from a specific host, group of hosts, or a subnet, and can also include source and destination TCP/UDP criteria. This allows you to apply a notify-only, throttling, or blocking policy while allowing exceptions for legitimate traffic from specific sources. You can also allow exceptions for traffic with specific TCP or UDP criteria.

For more information on when to apply connection-rate ACLs, refer to “Application Options” on page 3-6.

Note

Connection-rate ACLs are a special case of the switch’s ACL feature. If you need information on other applications of ACLs or more detailed information on how ACLs operate, refer to chapter 10, “Access Control Lists (ACLs)”.

Connection-Rate ACL Operation

A connection-rate ACL applies to inbound traffic on all ports configured for connection-rate filtering in the assigned VLAN, and creates an exception to the connection-rate filter policy configured on each port. A connection-rate ACL has no effect on ports in the VLAN that are not configured for connection-rate filtering.

A connection-rate ACL accepts inbound, legitimate traffic from trusted sources without filtering the traffic for the configured connection-rate policy. You can configure an ACL to assign policy filtering (**filter**) for traffic from some sources and no policy filtering (**ignore**) for traffic from other sources. However, the implicit **filter** invoked as the last entry in any connection-rate ACL ensures that any traffic not specifically excluded from policy filtering (by the **ignore** command) will be filtered by the configured policy for the port on which that traffic entered the switch.




Figure 3-8. Connection-Rate ACL Applied to Traffic Received Through a Given Port

Configuring a Connection-Rate ACL Using Source IP Address Criteria

(To configure a connection-rate ACL using UDP/TCP criteria, go to page 3-23.)

Syntax: ip access-list connection-rate-filter < crf-list-name >

Creates a connection-rate-filter ACL and puts the CLI into the access control entry (ACE) context:

```
ProCurve (config-crf-nacl) #
```

If the ACL already exists, this command simply puts the CLI into the ACE context.

Syntax: < filter | ignore > ip < any | host < ip-addr > | ip-addr < mask-length > >

Used in the ACE context (above) to specify the action of the connection-rate ACE and the source IP address of the traffic that the ACE affects.

< filter | ignore >

The **filter** option assigns policy filtering to traffic with source IP address (SA) matching the source address in the ACE. The **ignore** option specifies bypassing policy filtering for traffic with an SA that matches the source address in the ACE.

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

Specifies the SA criteria for traffic addressed by the ACE.

any: *Applies the ACEs action (**filter** or **ignore**) to traffic having any SA.*

host < ip-addr >: *Applies the ACEs action (**filter** or **ignore**) to traffic having the specified host SA.*

ip-addr < mask-length >: *Applies the ACEs action (**filter** or **ignore**) to traffic having an SA within the range defined by either:*

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

or

< src-ip-addr < mask >>

Use this criterion for traffic 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 ACE Mask” on page 3-26.

Configuring a Connection-Rate ACL Using UDP/TCP Criteria

(To configure a connection-rate ACL using source IP address criteria, refer to page 3-21.)

Syntax: ip access-list connection-rate-filter < crf-list-name >

Creates a connection-rate-filter ACL and puts the CLI into the access control entry (ACE) context:

```
ProCurve (config-crf-nacl) #
```

If the ACL already exists, this command simply puts the CLI into the ACE context.

Syntax: < filter | ignore > < udp | tcp > < any >

< filter | ignore > < udp | tcp > < host < ip-addr > > [udp/tcp-options]

< filter | ignore > < udp | tcp > < ip-addr < mask-length > > [udp/tcp-options]

Used in the ACE context (above) to specify the action of the connection-rate ACE (filter or ignore), and the UDP/TCP criteria and SA of the IP traffic that the ACE affects.

< filter | ignore >

filter: *This option assigns a policy of filtering (dropping) IP traffic having an SA that matches the source address criteria in the ACE.*

ignore: *This option specifies a policy of allowing IP traffic having an SA that matches the source address criteria in the ACE.*

< udp | tcp > < any | host < ip-addr > | ip-addr < mask-length > >

Applies the filter or ignore action to either TCP packets or UDP packets having the specified SA.

any: *Applies the ACEs action (filter or ignore) to IP traffic having any SA.*

host < ip-addr >: *Applies the ACEs action (filter or ignore) to IP traffic having the specified host SA.*

ip-addr < mask-length >: Applies the ACEs action (**filter** or **ignore**) to IP traffic having an SA within the range defined by either:

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

Use this criterion for traffic 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 ACE Mask” on page 3-26.

[udp/tcp-options]

destination-port < tcp-data > [source-port < tcp-data >]
source-port < tcp-data > [destination-port < tcp-data >]

destination-port < udp-data > [source-port < udp-data >]
source-port < udp-data > [destination-port < udp-data >]

tcp-data: < operator > < tcp-port-# >
udp-data: < operator > < udp-port-# >

operator: < eq | gt | lt | neq | range >

eq < port-nbr-or-name >: “Equal To”; to have a match with the ACE entry, the TCP or UDP source-port number in a packet must be equal to the specified port number.

gt: < port-nbr-or-name >: “Greater Than”; to have a match with the ACE entry, the TCP or UDP source-port number in a packet must be greater than the specified port number.

lt < port-nbr-or-name >: “Less Than”; to have a match with the ACE entry, the TCP or UDP source-port number in a packet must be less than the specified port number.

neq < port-nbr-or-name >: “Not Equal”; to have a match with the ACE entry, the TCP or UDP source-port number in a packet must not be equal to the specified port number.

range < start-port-nbr/name > < end-port-nbr/name >: To have a match with the ACE entry, the TCP or UDP source-port number in a packet must be in the range < start-port-nbr /name > < end-port-nbr/name >.

< tcp-data > or < udp-data >

TCP or UDP Port Number or (Well-Known) Port Name: Use the TCP or UDP port number required for the desired match. The switch also accepts certain well-known TCP or UDP port names as alternates to their corresponding port numbers:

TCP/UDP-PORT: Specify port by number.

bootpc: Bootstrap Protocol, client (68)

bootps: Bootstrap Protocol, server (67)

dns: Domain Name Service (53)

ntp: Network Time Protocol (123)

radius: Remote Authentication Dial-In User Service (1812)

radius-old: Remote Authentication Dial-In User Service 1645)

rip: Routing Information Protocol (520)

snmp: Simple Network Management Protocol (161)

snmp-trap: Simple Network Management Protocol (162)

tftp: Trivial File Transfer Protocol (69)

```
ProCurve(config)# ignore tcp host 15.75.10.11 destination-port eq 1812  
source-port eq 1812
```

Ignore (allow) tcp traffic from the host at 15.75.10.11 with both source and destination tcp ports of 1812.

```
ProCurve(config)# filter udp 15.75.10.0/24 source-port neq 162  
destination-port eq 162
```

Filter (drop) udp traffic from the subnet at 15.75.10.0 with a source udp port number not equal to 162 and a destination udp port number of 162.

Figure 3-9. Examples of Connection-Rate ACEs Using UDP/TCP Criteria

Applying Connection-Rate ACLs

To apply a connection-rate ACL, use the access group command described below. Note that this command differs from the access group command for non-connection-rate ACLs.

Syntax: [no] vlan < vid > ip access-group < crf-list-name > connection-rate-filter

*This command applies a connection-rate access control list (ACL) to inbound traffic on ports in the specified VLAN that are configured for connection-rate filtering. (A connection-rate ACL does not apply to ports in the VLAN that are not configured for connection-rate filtering.) The **no** form of the command removes the connection-rate ACL assignment from the VLAN.*

Note: *The switch allows only one connection-rate ACL assignment per VLAN. If a connection-rate ACL is already assigned to a VLAN and you assign another connection-rate ACL to that VLAN, the second ACL overwrites the first one. (A connection-rate ACL can be in addition to any standard or extended ACLs already assigned to the VLAN.)*

Using CIDR Notation To Enter the ACE Mask

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

Table 3-2. Examples of CIDR Notation for Masks

IP Address Used In an ACL with CIDR Notation	Resulting ACL Mask	Meaning
10.38.240.125/15	0.1.255.255	The leftmost 15 bits must match; the remaining bits are wildcards.
10.38.240.125/20	0.0.15.255	The leftmost 20 bits must match; the remaining bits are wildcards.
10.38.240.125/21	0.0.7.255	The leftmost 21 bits must match; the remaining bits are wildcards.
10.38.240.125/24	0.0.0.255	The leftmost 24 bits must match; the remaining bits are wildcards.
10.38.240.125/32	0.0.0.0	All bits must match.

For more on ACE masks, refer to “How an ACE Uses a Mask To Screen Packets for Matches” on page 10-36.

Example of Using an ACL in a Connection-Rate Configuration

This example adds connection-rate ACLs to the basic example on page 3-13.

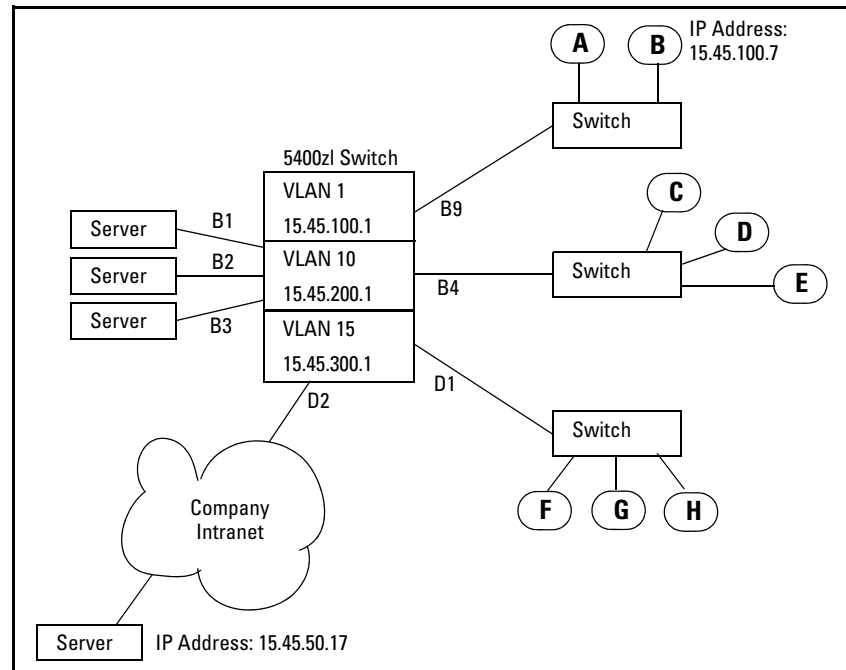


Figure 3-10. Sample Network

In the basic example on page 3-13, the administrator configured connection-rate blocking on port D2. However:

- The administrator has elevated the connection-rate sensitivity to **high**.
- The server at IP address 15.45.50.17 frequently transmits a relatively high rate of legitimate connection requests, which now triggers connection-rate blocking of the server’s IP address on port D2. This causes periodic, unnecessary blocking of access to the server.

The administrator needs to maintain blocking protection from the “Company Intranet” while allowing access to the server at 15.45.50.17. Because the server is carefully maintained as a trusted device, the administrator’s solution is to

configure a connection-rate ACL that causes the switch to ignore (circumvent) connection-rate filtering for inbound traffic from the server, while maintaining the filtering for all other inbound routed traffic on port D2.

The configuration steps include:

1. Create the connection-rate ACL with a single entry:
 - Use the IP address of the desired server.
 - Include a CIDR notation of “32” for the ACL mask. (Which means the mask will allow only traffic whose source IP address (SA) exactly matches the specified IP address.)
 - The ACL will automatically include the implicit **filter** ACE as the last entry, which means that any traffic that is not from the desired server will be subject to filtering by the connection-rate policy configured on port D2.
2. Assigning the ACL to the VLAN through which traffic from the server enters the switch.

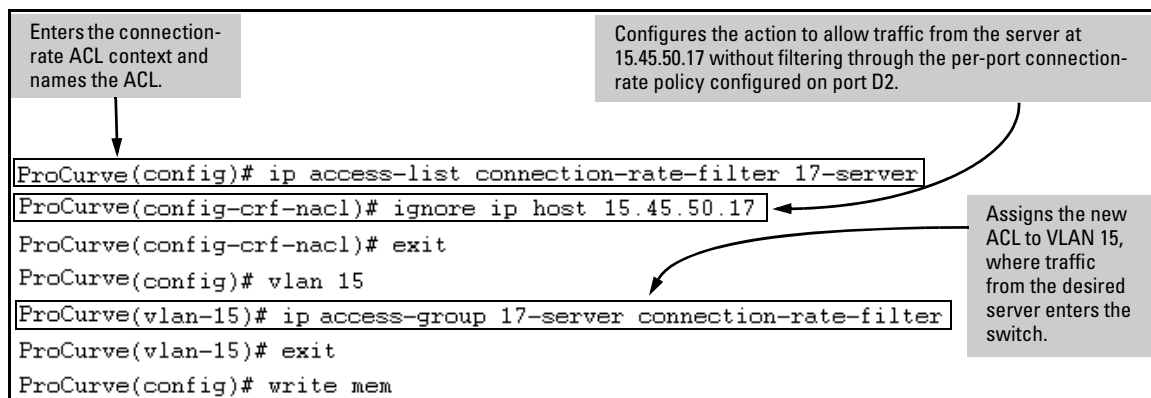


Figure 3-11. Creating and Assigning a Connection Rate ACL


```
ProCurve(config)# show config
Startup configuration:

; J8697A Configuration Editor; Created on release #K.11.XX

hostname "ProCurve"
connection-rate-filter sensitivity high
ip access-list connection-rate-filter "17-server"
  ignore ip 15.45.50.17 0.0.0.0
  exit
module 2 type J8702A
module 4 type J8702A
ip routing
snmp-server community "public" Unrestricted
snmp-server host 15.45.200.75 "public"
vlan 1
  name "DEFAULT_VLAN"
  untagged B5-B24
  ip address dhcp-bootp
  no untagged B1-B4,D1-D24
  ip proxy-arp
  exit
vlan 10
  name "VLAN10"
  untagged B1-B4
  no ip address
  ip proxy-arp
  exit
vlan 15
  name "VLAN15"
  untagged D1-D24
  no ip address
  ip proxy-arp
  ip access-group "17-server" connection-rate-filter
  exit
filter connection-rate B4 notify-only
filter connection-rate B1-B3 throttle
filter connection-rate B9.D1-D2 block
```

The new switch configuration includes the ACL configured in figure 3-11.

Shows the assignment of the above connection-rate ACL to VLAN 15.

Figure 3-12. Example of Switch Configuration Display with a Connection-Rate ACL

Connection-Rate ACL Operating Notes

- **ACE Types:** A connection-rate ACL allows you to configure two types of ACEs (Access Control Entries):
 - **ignore < source-criteria >:** This ACE type directs the switch to permit all inbound traffic meeting the configured < source-criteria > without filtering the traffic through the connection-rate policy configured on the port through which the traffic entered the switch. For example, **ignore host 15.45.120.70** tells the switch to permit traffic from the host at 15.45.120.70 without filtering this host's traffic through the connection-rate policy configured for the port on which the traffic entered the switch.

- **filter < source-criteria >**: This ACE type does the opposite of an **ignore** entry. That is, all inbound traffic meeting the configured **< source-criteria >** must be filtered through the connection-rate policy configured for the port on which the traffic entered the switch. This option is most useful in applications where it is easier to use **filter** to specify suspicious traffic sources for screening than to use **ignore** to specify exceptions for trusted traffic sources that don't need screening. For example, if the host at 15.45.127.43 requires connection-rate screening, but all other hosts in the VLAN do not, you would configure and apply a connection-rate ACL with **filter ip host 15.45.127.43** as the first ACE and **ignore ip any** as the second ACE. In this case, the traffic from host 15.45.127.43 would be screened, but traffic from all other hosts on the VLAN would be permitted without connection-rate screening.
- **Implicit ACE**: A connection-rate ACL includes a third, implicit **filter ip any** ACE which is automatically the last ACE in the ACL. This implicit ACE does not appear in displays of the ACL configuration, but is always present in any connection-rate ACL you configure. For example, assume that a port is configured with a connection-rate policy and is in a VLAN configured with a connection-rate ACL. If there is no match between an incoming packet and the ACE criteria in the ACL, then the implicit **filter ip any** sends the packet for screening by the connection-rate policy configured on that port. To preempt the implicit **filter ip any** in a given connection-rate ACL, you can configure **ignore IP any** as the last explicit ACE in the connection-rate ACL. The switch will then ignore (permit) traffic that is not explicitly addressed by other ACEs configured sequentially earlier in the ACL without filtering the traffic through the existing connection-rate policy.
- **Monitoring Shared Resources**: Active instances of throttling or blocking a client that is generating a high rate of connection requests uses internal routing switch resources that are shared with several other features. The routing switch provides ample resources for all features. However, if the internal resources become fully subscribed, new instances of throttling or blocking cannot be initiated until the necessary resources are released from other uses. (Event Log messages and SNMP traps are not affected.) For information on determining current resource availability and usage, refer to the appendix titled "Monitoring Resources" in the *Management and Configuration Guide* for your switch.

Connection-Rate Log and Trap Messages

These messages appear in the switch's Event Log identifying the source IP address of a connection-rate filtering event. If SNMP trap receivers are configured on the switch, it also sends the messages to the designated receiver(s).

Message	Meaning
Address not found in list of blocked hosts.	Appears in the CLI when the connection-rate-filter unblock command has been executed to unblock hosts that are not currently blocked.
W <mm/dd/yy hh:mm:ss> 00694 FFI: Src IP address <xxx.xxx.xxx.xxx> high connection rate, port <port number>	A warning that results when a port configured for notify-only detects a relatively high number of connection-rate attempts from a host.
W <mm/dd/yy hh:mm:ss> 00695 FFI: Src IP address <xxx.xxx.xxx.xxx> throttled, port <port number>	A warning and indication of the switch's response when a port configured for throttle detects a relatively high number of connection-rate attempts from a host.
W <mm/dd/yy hh:mm:ss> 00696 FFI: Src IP address <xxx.xxx.xxx.xxx> blocked, port <port number>	A warning and indication of the switch's response when a port configured for block detects a relatively high number of connection-rate attempts from a host.

— This page is intentionally unused —

Web and MAC Authentication

Contents

Overview	4-2
Client Options	4-3
General Features	4-3
How Web and MAC Authentication Operate	4-5
Authenticator Operation	4-5
Terminology	4-9
Operating Rules and Notes	4-10
General Setup Procedure for Web/MAC Authentication	4-12
Do These Steps Before You Configure Web/MAC Authentication ..	4-12
Additional Information for Configuring the RADIUS Server To Support MAC Authentication	4-13
Configuring the Switch To Access a RADIUS Server	4-14
Configuring Web Authentication on the Switch	4-17
Overview	4-17
Configure the Switch for Web-Based Authentication	4-18
Configuring MAC Authentication on the Switch	4-24
Overview	4-24
Configure the Switch for MAC-Based Authentication	4-25
Show Commands for Web-Based Authentication	4-28
Example: Verifying a Web Authentication Configuration	4-29
Configuring MAC Authentication	4-31
Configuration Overview	4-31
Config Commands for MAC-Based Authentication	4-31
Show Commands for MAC-Based Authentication	4-36
Example: Verifying a MAC Authentication Configuration	4-38
Client Status	4-39

Overview

Feature	Default	Menu	CLI	Web
Configure Web Authentication	n/a	—	4-17	—
Configure MAC Authentication	n/a	—	4-24	—
Display Web Authentication Status and Configuration	n/a	—	4-28	—
Display MAC Authentication Status and Configuration	n/a	—	4-36	—

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 32 clients allowed on the port. (The default is one client.)

Web authentication, MAC authentication, MAC lockdown, MAC lockout, and port-security are mutually exclusive on a given port. Also, LACP must be disabled on ports configured for any of these authentication methods.

Client Options

Web-Auth and MAC-Auth provide a port-based solution in which a port can belong to one, untagged VLAN at a time. However, where all clients can operate in the same VLAN, the switch allows up to 32 simultaneous clients per port. (In applications where you want the switch to simultaneously support multiple client sessions in different VLANs, design your system so that such clients will use different switch ports.)

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

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

General Features

Web and MAC Authentication on the 5400zl switches 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.

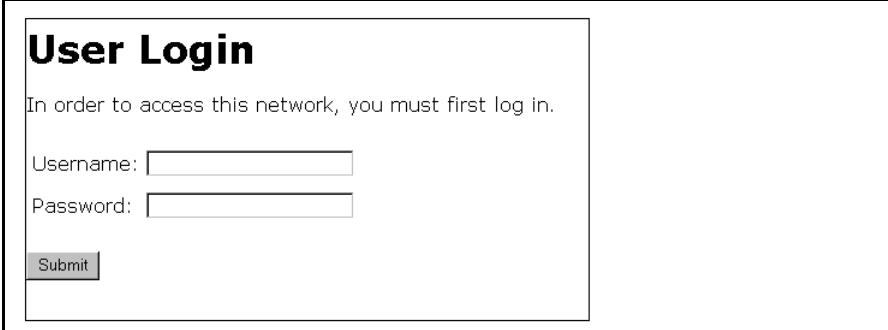
How Web and MAC Authentication Operate

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.



User Login

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

Username:

Password:

Figure 4-1. 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 4-2. 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 4-3. 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.

Web and MAC Authentication

How Web and MAC Authentication Operate

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

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

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

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

Terminology

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

Authentication Server: The entity providing an authentication service to the switch. In the case of a Series 5400zl switch running Web/MAC-Authentication, this is a RADIUS server.

Authenticator: In ProCurve switch applications, a device such as a Series 5400zl switch 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 4-21.

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

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

Operating Rules and Notes

- The switch supports concurrent 802.1X and either Web- or MAC-authentication operation on a port (with up to 32 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):
 - a. MAC lockout
 - b. MAC lockdown or Port Security
 - c. 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 the port 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 4-12.)
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 4-14.)

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:

Web and MAC Authentication

Configuring the Switch To Access a RADIUS Server

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

aabbccddeeff (the default format)

aabbcc-ddeeff

aa-bb-cc-dd-ee-ff

aa:bb:cc:dd:ee:ff

Note on MAC Addresses

Letters in MAC addresses must be in lowercase.

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

Configuring the Switch To Access a RADIUS Server

RADIUS Server Configuration Commands

radius-server	
[host <ip-address>]	below
[key <global-key-string >]	below
radius-server host <ip-address> key <server-specific key-string>	4-15

This section describes the minimal commands for configuring a RADIUS server to support Web-Auth and MAC Auth. For information on other RADIUS command options, refer to chapter 6, “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 6-1.)*

[key < global-key-string >]

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

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

Optional. Specifies an encryption key for use during authentication (or accounting) sessions with the 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.*

Web and MAC Authentication

Configuring the Switch To Access a RADIUS 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 '1A7rd'

```
ProCurve(config)# radius-server host 192.168.32 11
ProCurve(config)# radius-server host 192.168.32.11 key 1A7rd

ProCurve(config)# show radius

Status and Counters - General RADIUS Information

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

Server IP Addr   Auth  Acct
                Port  Port  Encryption Key
-----
192.168.32.11   1812 1813  1A7rd

ProCurve(config)#
```

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

Configuring Web Authentication on the Switch

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	4-18
aaa port-access web-based dhcp-lease	4-18
[no] aaa port-access web-based [e] <port-list>	4-19
[auth-vid]	4-19
[client-limit]	4-19
[client-moves]	4-19
[logoff-period]	4-20
[max-requests]	4-20
[max-retries]	4-20
[quiet-period]	4-20
[reauth-period]	4-20
[reauthenticate]	4-20
[redirect-url]	4-21
[server-timeout]	4-21
[ssl-login]	4-21
[unauth-vid]	4-27

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

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

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

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

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

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

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

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

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

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

Syntax: aaa port-access web-based [e] < port-list> [client-limit <1-32>]

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

Note: *On switches where Web Auth and 802.1X can operate concurrently, this limit includes the total number of clients authenticated through both methods.*

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Forces a reauthentication of all attached clients on the port.

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

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

Note: *The `redirect-url` command accepts only the first 103 characters of the allowed 127 characters.*

*Use the **no** form of the command to remove a specified redirect URL.*

(Default: There is no default URL. Browser behavior for authenticated clients may not be acceptable.)

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

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

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

Enables or disables SSL login (`https` on port 443). SSL must be enabled on the switch.

If SSL login is enabled, a user is redirected to a secure page, where they enter their username and password. If SSL login is disabled, a user is not redirected to a secure page to enter their credentials.

*Use the **no** form of the command to disable SSL login.
(Default: disabled)*

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

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

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

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

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

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

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

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

- *The 802.1s Multiple Spanning Tree Protocol (MSTP) or 802.1w Rapid Spanning Tree Protocol (RSTP) is enabled on the switch. MSTP and RSTP improve resource utilization while maintaining a loop-free network.*
- *The port is configured as an edge port in the network using the **spanning-tree edge-port** command.*

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

Notes:

- For information on how to configure the prerequisites for using the **aaa port-access controlled-directions in** command, see Chapter 4, “Multiple Instance Spanning-Tree Operation” in the Advanced Traffic Management Guide.
- To display the currently configured Controlled Directions value for web-authenticated ports, enter the **show port-access web-based config** command as shown in Figure 4-5.
- The **aaa port-access controlled-direction in** command allows Wake-on-LAN traffic to be transmitted on a web-authenticated egress port that has not yet transitioned to the authenticated state; the **controlled-direction both** setting prevents Wake-on-LAN traffic to be transmitted on a web-authenticated egress port until authentication occurs.

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

- Using the **aaa port-access controlled-directions in** command, you can enable the transmission of Wake-on-LAN traffic on unauthenticated egress ports that are configured for any of the following port-based security features:
 - 802.1X authentication
 - MAC 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 802.1X authentication, refer to chapter 13, “Configuring Port-Based and User-Based Access Control (802.1X)”.

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

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	4-25
[no] aaa port-access mac-based [e] < port-list >	4-25
[addr-limit]	4-26
[addr-moves]	4-26
[auth-vid]	4-26
[logoff-period]	4-26
[max-requests]	4-26
[quiet-period]	4-27
[reauth-period]	4-27
[reauthenticate]	4-27
[server-timeout]	4-27
[unauth-vid]	4-27

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

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

no-delimiter — specifies an aabbccddeeff format.

single-dash — specifies an aabbcc-ddeeff format.

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

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

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 [e] < port-list > [addr-limit <1-32>]

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

Note: *On switches where MAC Auth and 802.1X can operate concurrently, this limit includes the total number of clients authenticated through both methods.*

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

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

(Default: disabled – no moves allowed)

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

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

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

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

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

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

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

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

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

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

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

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

Forces a reauthentication of all attached clients on the port.

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

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

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

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

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

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

Show Commands for Web-Based Authentication

Command	Page
show port-access [<i>port-list</i>] web-based	4-28
[clients]	4-28
[config]	4-28
[config [auth-server]]	4-29
[config [web-server]]	4-29
show port-access <i>port-list</i> web-based config detail	4-29

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
Configuring MAC Authentication on the Switch

```
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

  Port  Enabled  Client  Client  Logoff  Re-Auth  Unauth  Auth  Cntrl
  ----  -
  1     No       1       No      300     0         0       0     both
  2     No       1       No      300     0         0       0     in
  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
  9     No       1       No      300     0         0       0     both
  10    No       1       No      300     0         0       0     both
  11    No       1       No      300     0         0       0     both
  12    No       1       No      300     0         0       0     both
  13    No       1       No      300     0         0       0     both
```

Figure 4-5. Example of Verifying a Web Authentication Configuration

Configuring MAC Authentication

Configuration 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 by configuring 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.

Config Commands for MAC-Based Authentication

Command	Page
aaa port-access mac-based addr-format	4-25
[no] aaa port-access mac-based [e] <port-list>	4-25
[addr-limit]	4-26
[addr-moves]	4-26
[auth-vid]	4-26
[logoff-period]	4-26
[max-requests]	4-26
[quiet-period]	4-27
[reauth-period]	4-27
[reauthenticate]	4-27
[server-timeout]	4-27
[unauth-vid]	4-27
aaa port-access <port-list> controlled-directions <both in>	4-34

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

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

no-delimiter — specifies an aabbccddeeff format.

single-dash — specifies an aabbcc-ddeeff format.

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

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

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

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

Syntax: aaa port-access mac-based [e] < port-list > [addr-limit <1-32>]

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

Note: *On switches where MAC Auth and 802.1X can operate concurrently, this limit includes the total number of clients authenticated through both methods.*

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

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

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

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

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

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

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

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

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

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

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

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

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

Forces a reauthentication of all attached clients on the port.

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

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

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

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

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

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

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

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

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

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

- *The 802.1s Multiple Spanning Tree Protocol (MSTP) or 802.1w Rapid Spanning Tree Protocol (RSTP) is enabled on the switch. MSTP and RSTP improve resource utilization while maintaining a loop-free network.*
- *The port is configured as an edge port in the network using the **spanning-tree edge-port** command.*

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

*To display the currently configured Controlled Directions value for MAC-authenticated ports, enter the **show port-access mac-based config** command as shown in Figure 4-6.*

Notes:

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

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

Notes: — Continued —

- 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 802.1X authentication, refer to chapter 13, “Configuring Port-Based and User-Based Access Control (802.1X)”.

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

Show Commands for MAC-Based Authentication

Command	Page
show port-access [<i>port-list</i>] mac-based	4-36
[clients]	4-37
[config]	4-37
[config [auth-server]]	4-37
show port-access <i>port-list</i> mac-based config detail	4-37

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

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             in
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
9      No        1       No            300           0               0               0             both
10     No        1       No            300           0               0               0             both
11     No        1       No            300           0               0               0             both
12     No        1       No            300           0               0               0             both
13     No        1       No            300           0               0               0             both
```

Figure 4-6. Example of Verifying a MAC Authentication Configuration

Client Status

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

Reported Status	Available Network Connection	Possible Explanations
authenticated	Authorized VLAN	Client authenticated. Remains connected until logoff-period or reauth-period expires.
authenticating	Switch only	Pending RADIUS request.
rejected-no vlan	No network access	<ol style="list-style-type: none"> 1. Invalid credentials supplied. 2. RADIUS Server difficulties. See log file. 3. If unauth-vid is specified it cannot be successfully applied to the port. An authorized client on the port has precedence.
rejected-unauth vlan	Unauthorized VLAN only	<ol style="list-style-type: none"> 1. Invalid credentials supplied. 2. RADIUS Server difficulties. See log file.
timed out-no vlan	No network access	RADIUS request timed out. If unauth-vid is specified it cannot be successfully applied to the port. An authorized client on the port has precedence. Credentials resubmitted after quiet-period expires.
timed out-unauth vlan	Unauthorized VLAN only	RADIUS request timed out. After the quiet-period expires credentials are resubmitted when client generates traffic.
unauthenticated	Switch only	Waiting for user credentials.

— This page is intentionally unused —

TACACS+ Authentication

Contents

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

Overview

Feature	Default	Menu	CLI	Web
view the switch's authentication configuration	n/a	—	page 5-9	—
view the switch's TACACS+ server contact configuration	n/a	—	page 5-10	—
configure the switch's authentication methods	disabled	—	page 5-11	—
configure the switch to contact TACACS+ server(s)	disabled	—	page 5-15	—

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

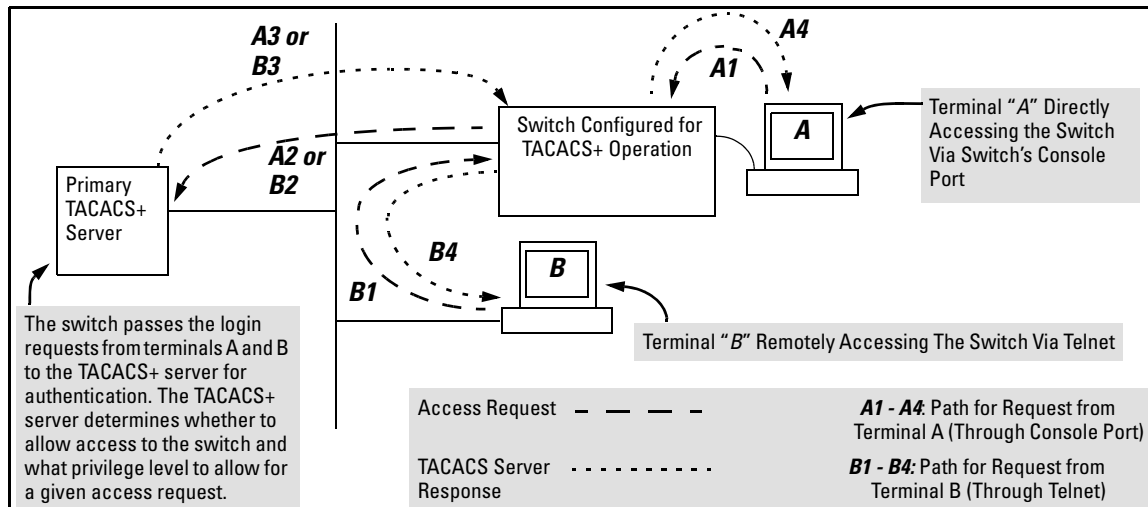


Figure 5-1. Example of TACACS+ Operation

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

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

TACACS+ does not affect web browser interface access. See “Controlling Web Browser Interface Access” on page 5-24.

Terminology Used in TACACS Applications:

- **NAS (Network Access Server):** This is an industry term for a TACACS-aware device that communicates with a TACACS server for authentication services. Some other terms you may see in literature describing TACACS operation are *communication server*, *remote access server*, or *terminal server*. These terms apply to a switch when TACACS+ is enabled on the switch (that is, when the switch is TACACS-aware).
- **TACACS+ Server:** The server or management station configured as an access control server for TACACS-enabled devices. To use TACACS+ with a switch covered in this guide and any other TACACS-capable devices in your network, you must purchase, install, and configure a TACACS+ server application on a networked server or management station in the network. The TACACS+ server application you install will provide various options for access control and access notifications. For more on the TACACS+ services available to you, see the documentation provided with the TACACS+ server application you will use.
- **Authentication:** The process for granting user access to a device through entry of a user name and password and comparison of this username/password pair with previously stored username/password data. Authentication also grants levels of access, depending on the privileges assigned to a user name and password pair by a system administrator.
 - **Local Authentication:** This method uses username/password pairs configured locally on the switch; one pair each for manager-level and operator-level access to the switch. You can assign local usernames and passwords through the CLI or web browser inter-

TACACS+ Authentication

Terminology Used in TACACS Applications:

face. (Using the menu interface you can assign a local password, but not a username.) Because this method assigns passwords to the switch instead of to individuals who access the switch, you must distribute the password information on each switch to everyone who needs to access the switch, and you must configure and manage password protection on a per-switch basis. (For more on local authentication, refer to chapter 2, “Configuring Username and Password Security”.)

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

General System Requirements

To use TACACS+ authentication, you need the following:

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

Notes

The effectiveness of TACACS+ security depends on correctly using your TACACS+ server application. For this reason, 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 5-24.

General Authentication Setup Procedure

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

other access type (console, in this case) open in case the Telnet access fails due to a configuration problem. The following procedure outlines a general setup procedure.

Note

If a complete access lockout occurs on the switch as a result of a TACACS+ configuration, see “Troubleshooting TACACS+ Operation” in the 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 the 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 5-23.)
2. Determine the following:
 - The IP address(es) of the TACACS+ server(s) you want the switch to use for authentication. If you will use more than one server, determine which server is your first-choice for authentication services.
 - The encryption key, if any, for allowing the switch to communicate with the server. You can use either a global key or a server-specific key, depending on the encryption configuration in the TACACS+ server(s).
 - The number of log-in attempts you will allow before closing a log-in session. (Default: 3)
 - The period you want the switch to wait for a reply to an authentication request before trying another server.
 - The username/password pairs you want the TACACS+ server to use for controlling access to the switch.
 - The privilege level you want for each username/password pair administered by the TACACS+ server for controlling access to the switch.
 - The username/password pairs you want to use for local authentication (one pair each for Operator and Manager levels).
3. Plan and enter the TACACS+ server configuration needed to support TACACS+ operation for Telnet access (login and enable) to the switch. This includes the username/password sets for logging in at the Operator (read-only) privilege level and the sets for logging in at the Manager (read/write) privilege level.

**Note on
Privilege Levels**

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

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

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

Caution

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

5. Using a terminal device connected to the switch’s console port, configure the switch for TACACS+ authentication *only* for **telnet login** access and **telnet enable** access. At this stage, do not configure TACACS+ authentication for console access to the switch, as you may need to use the console for access if the configuration for the Telnet method needs debugging.
6. Ensure that the switch is configured to operate on your network and can communicate with your first-choice TACACS+ server. (At a minimum, this requires IP addressing and a successful **ping** test from the switch to the server.)
7. On a remote terminal device, use Telnet to attempt to access the switch. If the attempt fails, use the console access to check the TACACS+ configuration on the switch. If you make changes in the switch configuration, check Telnet access again. If Telnet access still fails, check the

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

8. After your testing shows that Telnet access using the TACACS+ server is working properly, configure your TACACS+ server application for console access. Then test the console access. If access problems occur, check for and correct any problems in the switch configuration, and then test console access again. If problems persist, check your TACACS+ server application for mis-configurations or missing data that could affect the console access.
9. When you are confident that TACACS+ access through both Telnet and the switch's console operates properly, use the **write memory** command to save the switch's running-config file to flash.

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 5-5 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	5-9
show tacacs	5-10
aaa authentication	5-11 through 5-14
console	
Telnet	
num-attempts <1-10 >	
tacacs-server	5-15
host < ip-addr >	5-15
key	5-19
timeout < 1-255 >	5-20

Viewing the Switch's Current Authentication Configuration

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

Syntax: show authentication

This example shows the default authentication configuration.

```

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

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

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

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

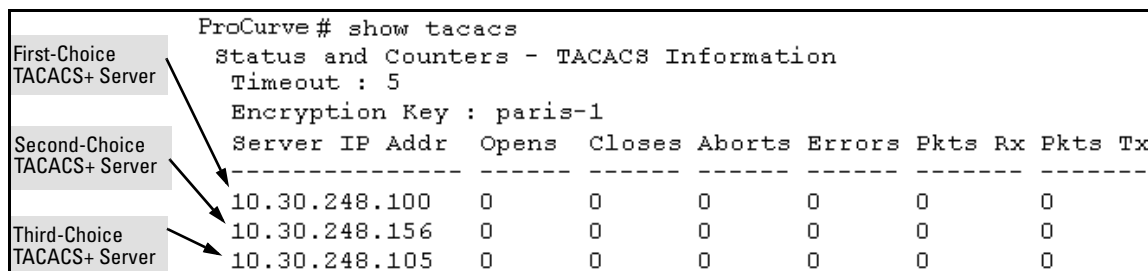
Figure 5-2. Example Listing of the Switch's Authentication Configuration

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 5-3. Example of the Switch's TACACS+ Configuration Listing

Configuring the Switch's Authentication Methods

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

Syntax: aaa authentication

< console | telnet >

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

< enable | login >

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

< local | tacacs | radius >

Selects the type of security access:

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

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

radius — Authenticates with a password and other data configured on a RADIUS server. (Refer to chapter 6, "RADIUS Authentication and Accounting".)

[< local | none >]

If the primary authentication method fails, determines whether to use the local password as a secondary method or to disallow access. Local is only available as a secondary method if the primary method is tacacs or radius.

aaa authentication num-attempts < 1-10 >

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

Table 5-1. AAA Authentication Parameters

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

As shown in the next table, login and enable access is always available locally through a direct terminal connection to the switch's console port. However, for Telnet access, you can configure TACACS+ to deny access if a TACACS+ server goes down or otherwise becomes unavailable to the switch.

Table 5-2. Primary/Secondary Authentication Table

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

Caution Regarding the Use of Local for Login Primary Access

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

TACACS+ Authentication

Configuring TACACS+ on the Switch

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

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

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

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

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

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

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

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

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

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

```
ProCurve (config)# aaa authentication num-attempts 2
```

Configuring the Switch's TACACS+ Server Access

The `tacacs-server` command configures these parameters:

- **The host IP address(es)** for up to three TACACS+ servers; one first-choice and up to two backups. Designating backup servers provides for a continuation of authentication services in case the switch is unable to contact the first-choice server.
- **An optional encryption key.** This key helps to improve security, and must match the encryption key used in your TACACS+ server application. In some applications, the term “secret key” or “secret” may be used instead of “encryption key”. If you need only one encryption key for the switch to use in all attempts to authenticate through a TACACS+ server, configure a global key. However, if the switch is configured to access multiple TACACS+ servers having different encryption keys, you can configure the switch to use different encryption keys for different TACACS+ servers.
- **The timeout value** in seconds for attempts to contact a TACACS+ server. If the switch sends an authentication request, but does not receive a response within the period specified by the timeout value, the switch resends the request to the next server in its Server IP Address list, if any. If the switch still fails to receive a response from any TACACS+ server, it reverts to whatever secondary authentication method was configured using the **aaa authentication** command (local or none; see “Configuring the Switch's Authentication Methods” on page 5-11.)

Note

As described under “General Authentication Setup Procedure” on page 5-5, 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”.

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

Specifies the IP address of a device running a TACACS+ server application. Optionally, can also specify the unique, per-server encryption key to use when each assigned server has its own, unique key. For more on the encryption key, see **"Using the Encryption Key" on page 5-23** and the documentation provided with your TACACS+ server application.

You can enter up to three IP addresses; one first-choice and two (optional) backups (one second-choice and one third-choice).

Use **show tacacs** to view the current IP address list.

If the first-choice TACACS+ server fails to respond to a request, the switch tries the second address, if any, in the show tacacs list. If the second address also fails, then the switch tries the third address, if any.

(See figure 5-3, "Example of the Switch's TACACS+ Configuration Listing" on 5-10.)

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 5-20.

TACACS+ Authentication
 Configuring TACACS+ on the Switch

Name	Default	Range
key <key-string>	none (null)	n/a
<p>Specifies the optional, global “encryption key” that is also assigned in the TACACS+ server(s) that the switch will access for authentication. This option is subordinate to any “per-server” encryption keys you assign, and applies only to accessing TACACS+ servers for which you have not given the switch a “per-server” key. (See the host <ip-addr> [key <key-string> entry at the beginning of this table.)</p> <p>For more on the encryption key, see “Using the Encryption Key” on page 5-23 and the documentation provided with your TACACS+ server application.</p>		
timeout <1 - 255>	5 sec	1 - 255 sec
<p>Specifies how long the switch waits for a TACACS+ server to respond to an authentication request. If the switch does not detect a response within the timeout period, it initiates a new request to the next TACACS+ server in the list. If all TACACS+ servers in the list fail to respond within the timeout period, the switch uses either local authentication (if configured) or denies access (if none configured for local authentication).</p>		

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

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

```

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

Figure 5-4. Example of the Switch with Two TACACS+ Server Addresses Configured

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

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

```
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
```

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

Figure 5-5. Example of the Switch After Assigning a Different "First-Choice" Server

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

```
ProCurve(config)# no tacacs-server host 10.28.227.15
```

Configuring an Encryption Key. Use an encryption key in the switch if the switch will be requesting authentication from a TACACS+ server that also uses an encryption key. (If the server expects a key, but the switch either does not provide one, or provides an incorrect key, then the authentication attempt will fail.) Use a *global encryption key* if the same key applies to all TACACS+ servers the switch may use for authentication attempts. Use a *per-server encryption key* if different servers the switch may use will have different keys. (For more details on encryption keys, see "Using the Encryption Key" on page 5-23.)

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

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

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

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

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

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

```
ProCurve(config)# no tacacs-server key
```

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

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

Note

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

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

```
ProCurve(config)# tacacs-server timeout 3
```

How Authentication Operates

General Authentication Process Using a TACACS+ Server

Authentication through a TACACS+ server operates generally as described below. For specific operating details, refer to the documentation you received with your TACACS+ server application.

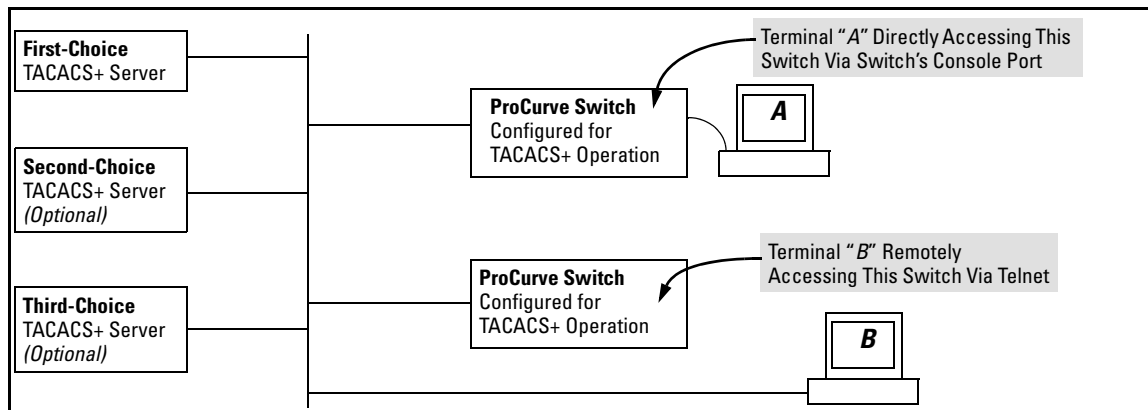


Figure 5-6. Using a TACACS+ Server for Authentication

Using figure 5-6, above, after either switch detects an operator's logon request from a remote or directly connected terminal, the following events occur:

1. The switch queries the first-choice TACACS+ server for authentication of the request.
 - If the switch does not receive a response from the first-choice TACACS+ server, it attempts to query a secondary server. If the switch does not receive a response from any TACACS+ server, then it uses its own local username/password pairs to authenticate the logon request. (See "Local Authentication Process" on page 5-22.)
 - If a TACACS+ server recognizes the switch, it forwards a username prompt to the requesting terminal via the switch.
2. When the requesting terminal responds to the prompt with a username, the switch forwards it to the TACACS+ server.
3. After the server receives the username input, the requesting terminal receives a password prompt from the server via the switch.
4. When the requesting terminal responds to the prompt with a password, the switch forwards it to the TACACS+ server and one of the following actions occurs:
 - If the username/password pair received from the requesting terminal matches a username/password pair previously stored in the server, then the server passes access permission through the switch to the terminal.
 - If the username/password pair entered at the requesting terminal does not match a username/password pair previously stored in the server, access is denied. In this case, the terminal is again prompted to enter a username and repeat steps 2 through 4. In the default configuration, the switch allows up to three attempts to authenticate a login session. If the requesting terminal exhausts the attempt limit without a successful TACACS+ authentication, the login session is terminated and the operator at the requesting terminal must initiate a new session before trying again.

Local Authentication Process

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

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

(For a listing of authentication options, see table 5-2, “Primary/Secondary Authentication Table” on 5-13.)

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

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

Note

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

Using the Encryption Key

General Operation

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

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

Note

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

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

Encryption Options in the Switch

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

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

```
ProCurve(config)# tacacs-server key north40campus
```

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

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

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

Controlling Web Browser Interface Access When Using TACACS+ Authentication

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

- Configure local authentication (a Manager user name and password and, optionally, an Operator user name and password) on the switch.
- Configure the switch's Authorized IP Manager feature to allow web browser access only from authorized management stations. (The Authorized IP Manager feature does not interfere with TACACS+ operation.)
- Disable web browser access to the switch by going to the System Information screen in the Menu interface and configuring the **Web Agent Enabled** parameter to **No**.

Messages Related to TACACS+ Operation

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

CLI Message	Meaning
Connecting to Tacacs server	The switch is attempting to contact the TACACS+ server identified in the switch's tacacs-server configuration as the first-choice (or only) TACACS+ server.
Connecting to secondary Tacacs server	The switch was not able to contact the first-choice TACACS+ server, and is now attempting to contact the next (secondary) TACACS+ server identified in the switch's tacacs-server configuration.
Invalid password	The system does not recognize the username or the password or both. Depending on the authentication method (tacacs or local), either the TACACS+ server application did not recognize the username/password pair or the username/password pair did not match the username/password pair configured in the switch.
No Tacacs servers responding	The switch has not been able to contact any designated TACACS+ servers. If this message is followed by the Username prompt, the switch is attempting local authentication.
Not legal combination of authentication methods	For console access , if you select tacacs as the primary authentication method, you must select local as the secondary authentication method. This prevents you from being locked out of the switch if all designated TACACS+ servers are inaccessible to the switch.
Record already exists	When resulting from a tacacs-server host <ip addr> command, indicates an attempt to enter a duplicate TACACS+ server IP address.

Operating Notes

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

- When TACACS+ is not enabled on the switch—or when the switch's only designated TACACS+ servers are not accessible—setting a local Operator password without also setting a local Manager password does not protect the switch from manager-level access by unauthorized persons.
- When using the **copy** command to transfer a configuration to a TFTP server, any optional, server-specific and global encryption keys (page 5-15) in the TACACS configuration will not be included in the transferred file. Otherwise, a security breach could occur, allowing access to the TACACS+ username/password information.

RADIUS Authentication and Accounting

Contents

Overview	6-3
Authentication Services	6-3
Accounting Services	6-4
RADIUS-Administered CoS and Rate-Limiting	6-4
SNMP Access to the Switch's Authentication Configuration MIB ...	6-4
Terminology	6-5
Switch Operating Rules for RADIUS	6-6
General RADIUS Setup Procedure	6-7
Configuring the Switch for RADIUS Authentication	6-8
Outline of the Steps for Configuring RADIUS Authentication	6-9
1. Configure Authentication for the Access Methods You Want RADIUS To Protect	6-10
2. Enable the (Optional) Access Privilege Option	6-12
3. Configure the Switch To Access a RADIUS Server	6-13
4. Configure the Switch's Global RADIUS Parameters	6-15
Using SNMP To View and Configure Switch Authentication Features 6-19	
Changing and Viewing the SNMP Access Configuration	6-20
Local Authentication Process	6-22
Controlling Web Browser Interface Access	6-23
Configuring RADIUS Authorization	6-24
Overview	6-24
Commands Authorization Type	6-24
Enabling Authorization with the CLI	6-25
Showing Authorization Information	6-26
Configuring the RADIUS Server	6-26
Using Vendor Specific Attributes (VSAs)	6-26

Example Configuration on Cisco Secure ACS for MS Windows	6-28
Example Configuration Using FreeRADIUS	6-30
Configuring RADIUS Accounting	6-32
Operating Rules for RADIUS Accounting	6-33
Steps for Configuring RADIUS Accounting	6-34
1. Configure the Switch To Access a RADIUS Server	6-35
2. Configure Accounting Types and the Controls for Sending Reports to the RADIUS Server	6-36
3. (Optional) Configure Session Blocking and Interim Updating Options	6-38
Viewing RADIUS Statistics	6-40
General RADIUS Statistics	6-40
RADIUS Authentication Statistics	6-42
RADIUS Accounting Statistics	6-43
Changing RADIUS-Server Access Order	6-44
Messages Related to RADIUS Operation	6-47

Overview

Feature	Default	Menu	CLI	Web
Configuring RADIUS Authentication	None	n/a	6-8	n/a
Configuring RADIUS Accounting	None	n/a	6-32	n/a
Configuring RADIUS Authorization	None	n/a	6-24	n/a
Viewing RADIUS Statistics	n/a	n/a	6-40	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 Services

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
- SFTP/SCP
- Web (5400zl, 4200vl, 2800s as of software version I.08.60, and 2600s as of software version H.08.58 switches)
- Port-Access (802.1X)

The switch also supports RADIUS accounting for Web Authentication and MAC authentication sessions.

Note

The switch does not support RADIUS security for SNMP (network management) access. For information on blocking access through the web browser interface, refer to “Controlling Web Browser Interface Access” on page 6-23.

Accounting Services

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.

RADIUS-Administered CoS and Rate-Limiting

The switches covered in this guide take advantage of vendor-specific attributes (VSAs) applied in a RADIUS server to support these optional, RADIUS-assigned attributes:

- 802.1p (CoS) priority assignment to inbound traffic on the specified port(s) (port-access authentication only)
- Per-Port Rate-Limiting on a port with an active link to an authenticated client (port-access authentication only)

SNMP Access to the Switch’s Authentication Configuration MIB

Beginning with software release K.12.*xx*, the switch’s default configuration allows SNMP access to the hpSwitchAuth MIB (Management Information Base). A management station running an SNMP networked device management application such as ProCurve Manager Plus (PCM+) or HP OpenView can access the switch’s MIB for read access to the switch’s status and read/write access to the switch’s configuration. For more information, including the CLI command to use for disabling this feature, refer to “Using SNMP To View and Configure Switch Authentication Features” on page 6-19.

Terminology

AAA: Authentication, Authorization, and Accounting groups of services provided by the carrying protocol.

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.

CoS (Class of Service): Support for priority handling of packets traversing the switch, based on the IEEE 802.1p priority carried by each packet. (For more on this topic, refer to the “Overview” section in the “Quality of Service (QoS)” chapter in the *Advanced Traffic Management Guide* for your switch.)

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).

EXEC Session: a service (EXEC shell) granted to the authenticated login user for doing management operations on the ProCurve device.

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): a protocol for carrying authentication, authorization, and accounting information between a Network Access Server and shared AAA servers in a distributed dial-in networking environment.

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.

Vendor-Specific Attribute: A vendor-defined value configured in a RADIUS server to specific an optional switch feature assigned by the server during an authenticated client session.

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 6-40). 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 6-44.)
- 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.
- When primary/secondary authentication is set to Radius/Local (for either Login or Enable) and the RADIUS server fails to respond to a client attempt to authenticate, the failure is noted in the Event Log with the message **radius: Can't reach RADIUS server < server-ip-addr >**. When this type of failure occurs, the switch prompts the client again to enter a username and password. In this case, use the local username (if any) and password configured on the switch itself.
- Zero-length usernames or passwords are not allowed for RADIUS authentication, even though allowed by some RADIUS servers.
- TACACS+ is not supported for the web browser interface access.

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 6-1. Preparation for Configuring RADIUS on the Switch

- Determine the access methods (console, Telnet, Port-Access (802.1X), web browser interface and/or SSH) for which you want RADIUS as the primary authentication method. Consider both Operator (login) and Manager (enable) levels, as well as which secondary authentication methods to use (local or none) if the RADIUS authentication fails or does not respond.

ProCurve(config)# show authentication					Note: The Webui access task shown in this figure is available only on the switches covered in this guide.
Status and Counters - Authentication Information					
Login Attempts : 3					
Respect Privilege : Disabled					
Access Task	Login Primary	Login Secondary	Enable Primary	Enable Secondary	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.
Console	Radius	Local	Radius	Local	
Telnet	Radius	Local	Radius	Local	
Port-Access	EapRadius				
Webui	Radius	Local	Radius	Local	
SSH	Radius	Local	Radius	Local	
Web-Auth	ChapRadius				
MAC-Auth	ChapRadius				

Figure 6-1. 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).

RADIUS Authentication and Accounting

Configuring the Switch for RADIUS Authentication

- 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.
 - Optional: Determine whether the switch access level (Manager or Operator) for authenticated clients can be set by a Service Type value the RADIUS server includes in its authentication message to the switch. (Refer to “2. Enable the (Optional) Access Privilege Option” on page 6-12.)
 - Configure RADIUS on the server(s) used to support authentication on the switch.
-

Configuring the Switch for RADIUS Authentication

RADIUS Authentication Commands	Page
aaa authentication	6-10
< console telnet ssh web > < enable login > radius*	6-10
[local none]	6-10
[login privilege-mode]*	6-12
[no] radius-server host < IP-address >	6-13
[auth-port < port-number >]	6-13
[acct-port < port-number >]	6-13, 6-35
[key < server-specific key-string >]	6-13
[no] radius-server key < global key-string >	6-16
radius-server timeout < 1 - 15 >	6-16
radius-server retransmit < 1 - 5 >	6-16
[no] radius-server dead-time < 1 - 1440 >	6-17
show radius	6-40
[< host < ip-address >]	6-41
show authentication	6-42
show radius authentication	6-43

*The **web** authentication option for the web browser interface is available on the switches covered in this guide.

Outline of the Steps for Configuring RADIUS Authentication

There are three main steps to configuring RADIUS authentication:

1. Configure RADIUS authentication for controlling access through one or more of the following
 - Serial port
 - Telnet
 - SSH
 - Port-Access (802.1X)
 - Web browser interface
2. Enable RADIUS authentication on the switch to override the default authentication operation of automatically assigning an authenticated client to the Operator privilege level. This optional feature applies the privilege level specified by the Service Type value received from the RADIUS server. (Refer to “1. Configure Authentication for the Access Methods You Want RADIUS To Protect” on page 6-10.)
3. 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)
4. Configure the global RADIUS parameters.
 - **Server Key:** This key must match the encryption key used on the RADIUS servers the switch contacts for authentication and accounting services unless you configure one or more per-server keys. (Default: null.)

- **Timeout Period:** The timeout period the switch waits for a RADIUS server to reply. (Default: 5 seconds; range: 1 to 15 seconds.)
- **Retransmit Attempts:** The number of retries when there is no server response to a RADIUS authentication request. (Default: 3; range of 1 to 5.)
- **Server Dead-Time:** The period during which the switch will not send new authentication requests to a RADIUS server that has failed to respond to a previous request. This avoids a wait for a request to time out on a server that is unavailable. If you want to use this feature, select a dead-time period of 1 to 1440 minutes. (Default: 0—disabled; range: 1 - 1440 minutes.) If your first-choice server was initially unavailable, but then becomes available before the dead-time expires, you can nullify the dead-time by resetting it to zero and then trying to log on again. As an alternative, you can reboot the switch, (thus resetting the dead-time counter to assume the server is available) and then try to log on again.
- **Number of Login Attempts:** This is actually an **aaa authentication** command. It controls how many times per session a RADIUS client (and 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 6-32.)

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 use RADIUS for SSH access, first configure the switch for SSH operation. Refer to chapter 8, “Configuring Secure Shell (SSH)” .
- **Web:** You can enable RADIUS authentication for web browser interface access to the switch.

You can configure RADIUS as the primary password authentication method for the above access methods. You also need to select either **local** or **none** as a secondary, or backup, method. Note that for console access, if you configure

radius (or **tacacs**) for primary authentication, you must configure **local** for the secondary method. This prevents the possibility of being completely locked out of the switch in the event that all primary access methods fail.

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

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

`[< local | none >]`

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

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

```

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

```

Status and Counters - Authentication Information

Login Attempts : 3
Respect Privilege : Disabled

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

Note: The Webui access task shown in this figure is available only on the switches covered in this guide.

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

Figure 6-2. Example Configuration for RADIUS Authentication

Note

If you configure the Login Primary method as **local** instead of **radius** (and local passwords are configured on the switch), then clients connected to your network 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 6-22.

2. Enable the (Optional) Access Privilege Option

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

Syntax: [no] aaa authentication login privilege-mode

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

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

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

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

— Continued on the next page. —

— Continued from the preceding page. —

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

3. Configure the Switch To Access a RADIUS Server

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

Note

If you want to configure RADIUS accounting on the switch, go to page 6-32: “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 6-44.)*

[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.

Note: When you save the config file using Xmodem or TFTP, the key information is not saved in the file. This causes Radius authentication to break when the config file is loaded back onto the switch.

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 6-3 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
  Port           Port   Port   Encryption Key
  -----
  10.33.18.127   1812  1813   TempKey01
```

Figure 6-3. Sample Configuration for RADIUS Server Before Changing the Key and Adding Another Server

To make the changes listed prior to figure 6-3, 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
  Port           Port   Port   Encryption Key
  -----
  10.33.18.127   1812  1813   source0127
  10.33.18.119   1812  1813   source0119
```

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

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

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

Figure 6-4. 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 6-44.

4. Configure the Switch's Global RADIUS Parameters

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

- **Number of login attempts:** In a given session, specifies how many tries at entering the correct username and password pair are allowed before access is denied and the session terminated. (This is a general **aaa authentication** parameter and is not specific to RADIUS.)

- **Global server key:** The server key the switch will use for contacts with all RADIUS servers for which there is not a server-specific key configured by **radius-server host < ip-address > key < key-string >**. This key is optional if you configure a server-specific key for each RADIUS server entered in the switch. (Refer to “3. Configure the Switch To Access a RADIUS Server” on page 6-13.)
- **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 6-5. Example of Global Configuration Exercise for RADIUS Authentication

RADIUS Authentication and Accounting

Configuring the Switch for RADIUS Authentication

```

ProCurve(config)# show authentication

Status and Counters - Authentication Information

Login Attempts : 2
Respect Privilege : Disabled
  
```

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

```

ProCurve(config)# show radius

Status and Counters - General RADIUS Information

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

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	

Note: The **Webui** access task shown in this figure is available only on the switches covered in this guide.

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

Global RADIUS parameters from figure 6-5.

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

Using SNMP To View and Configure Switch Authentication Features

Beginning with software release K.12.xxx, SNMP MIB object access is available for switch authentication configuration (hpSwitchAuth) features. This means that the switches covered by this Guide allow, by default, manager-only SNMP read/write access to a subset of the authentication MIB objects for the following features:

- number of primary and secondary login and enable attempts
- TACACS+ server configuration and status
- RADIUS server configuration
- selected 802.1X settings
- key management subsystem chain configuration
- key management subsystem key configuration
- OSPF interface authentication configuration
- local switch operator and manager usernames and passwords

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

Security Notes

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

To help prevent unauthorized access to the switch's authentication MIB, ProCurve recommends enhancing security according to the guidelines under "Switch Access Security" on page 1-3.

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

If you choose to leave SNMP access to the security MIB open (the default setting), ProCurve recommends that you configure the switch with the SNMP version 3 management and access security feature, and disable SNMP version

2c access. (Refer to “Switch Access Security” on page 1-3.)

Changing and Viewing the SNMP Access Configuration

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

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

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

Syntax: show snmp-server

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

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

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

SNMP Communities

Community Name	MIB View	Write Access
public	Manager	Unrestricted

Trap Receivers

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

Send Authentication Traps [No] : No

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

Excluded MIBs

```
hpSwitchAuthenticationMIB
```

This command disables SNMP security MIB access.

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

Figure 6-7. Disabling SNMP Access to the Authentication MIB and Displaying the Result

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

```
ProCurve(config)# show run

Running configuration:

; J8697A Configuration Editor; Created on release #K.12.01

hostname "ProCurve"
snmp-server mib hpSwitchAuthMIB excluded ] ← Indicates that SNMP access to the
ip default-gateway 10.10.24.55                authentication configuration MIB
snmp-server community "public" Operator      (hpSwitchAuth) is disabled.
vlan 1
  name "DEFAULT_VLAN"
  untagged A1-A24,B1-B4
  ip address 10.10.24.100 255.255.255.0
  exit
password manager
```

Figure 6-8. Using the show run Command to View the Current Authentication MIB Access State

Local Authentication Process

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

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

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

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

Controlling Web Browser Interface Access

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

- Configure the switch to support RADIUS authentication for web browser interface access (Web Authentication, Chapter 7).
- Options for the switches covered in this guide:
 - Configure local authentication (a Manager user name and password and, optionally, an Operator user name and password) on the switch.
 - Configure the switch's Authorized IP Manager feature to allow web browser access only from authorized management stations. (The Authorized IP Manager feature does not interfere with TACACS+ operation.)
 - Use one of the following methods to disable web browser access to the switch via http (Port 80):

CLI: **no web-management**

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

2. Switch Configuration

1. System Information

Web Agent Enabled: No

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 6-26.

After the Access-Accept packet is delivered, 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 6-9. 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\CiscoAAA\3.2\
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 >	6-35
[acct-port < port-number >]	6-35
[key < key-string >]	6-35
[no] aaa accounting < exec network system commands > < start-stop stop-only > radius	6-38
[no] aaa accounting update periodic < 1 - 525600 > (in minutes)	6-38
[no] aaa accounting suppress null-username	6-38
show accounting	6-43
show accounting sessions	6-44
show radius accounting	6-44

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 6-7 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 switches covered in this guide support four 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

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

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

Operating Rules for RADIUS Accounting

- You can configure up to four types of accounting to run simultaneously: exec, system, network, and commands.
- 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 6-44.)

- If access to a RADIUS server fails during a session, but after the client has been authenticated, the switch continues to assume the server is available to receive accounting data. Thus, if server access fails during a session, it will not receive accounting data transmitted from the switch.

Steps for Configuring RADIUS Accounting

1. Configure the switch for accessing a RADIUS server.

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

- Use the same **radius-server host** command that you would use to configure RADIUS authentication. Refer to “3. Configure the Switch To Access a RADIUS Server” on page 6-13.
 - 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 6-13. (Default: null)
2. Configure accounting types and the controls for sending reports to the RADIUS server.
 - **Accounting types:** exec (page 6-33), network (page 6-32), commands (page 6-33), or system (page 6-33)
 - **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 6-13. 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.

Note: When you save the config file using Xmodem or TFTP, the key information is not saved in the file. This causes Radius authentication to fail when the config file is loaded back onto the switch.

(For a more complete description of the **radius-server** command and its options, turn to page 6-13.)

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.).

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

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

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

Figure 6-10. Example of Configuring for a RADIUS Server with a Non-Default Accounting UDP Port Number

The radius-server command as shown in figure 6-10, 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 Services” on page 6-4.)
- **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 Services” on page 4.)
- **Commands:** When commands authorization is enabled, a record accounting notice is sent after the execution of a command.
- **Web or MAC:** You can also use Web or MAC to collect accounting information.

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, Commands, or System).
- Do not wait for an acknowledgement.

The system option (page 6-36) 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, Commands, or System).
- Do not wait for an acknowledgment.

The system option (page 6-36) always delivers **stop-only** operation because the switch sends the accumulated data only when there is a reboot, reload, or accounting on/off event.

Syntax: [no] aaa accounting < exec | network | system | commands > < start-stop | stop-only > radius

Configures RADIUS accounting type and how data will be sent to the RADIUS server.

For example, to configure RADIUS accounting on the switch with **start-stop** for exec functions and **stop-only** for system functions:

```
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
```

Configures exec and system accounting and controls.

Summarizes the switch's accounting configuration.

Exec and System accounting are active. (Assumes the switch is configured to access a reachable

Figure 6-11. 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 6-11, suppose that you wanted the switch to:

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

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

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

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

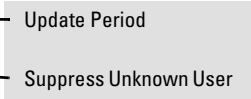


Figure 6-12. Example of Optional Accounting Update Period and Accounting Suppression on Unknown User

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 6-32.)*

```
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 6-13. Example of General RADIUS Information from Show Radius Command

```
ProCurve(config)# show radius host 192.33.12.65
Status and Counters - RADIUS Server Information
  Server IP Addr : 192.33.12.65
  Authentication UDP Port : 1812           Accounting UDP Port : 1813
  Round Trip Time      : 2                 Round Trip Time      : 7
  Pending Requests     : 0                 Pending Requests     : 0
  Retransmissions      : 0                 Retransmissions      : 0
  Timeouts             : 0                 Timeouts             : 0
  Malformed Responses  : 0                 Malformed Responses  : 0
  Bad Authenticators   : 0                 Bad Authenticators   : 0
  Unknown Types        : 0                 Unknown Types        : 0
  Packets Dropped      : 0                 Packets Dropped      : 0
  Access Requests      : 2                 Accounting Requests  : 2
  Access Challenges    : 0                 Accounting Responses : 2
  Access Accepts       : 2
  Access Rejects       : 0
```

Figure 6-14. RADIUS Server Information From the Show Radius Host Command

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

RADIUS Authentication 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 6-32.)*

ProCurve(config)# show authentication					Note: The Webui access task shown in this figure is available only on the 5400zl switches.
Status and Counters - Authentication Information					
Login Attempts : 2					
Respect Privilege : Disabled					
Access Task	Login Primary	Login Secondary	Enable Primary	Enable Secondary	
Console	Local	None	Local	None	
Telnet	Radius	None	Radius	None	
Port-Access	Local				
Webui	Local	None	Local	None	
SSH	Radius	None	Radius	None	
Web-Auth	ChapRadius				
MAC-Auth	ChapRadius				

Figure 6-15. 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 : ProCurve.
Invalid Server Addresses : 0

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

Figure 6-16. 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.

```

HPswitch # show accounting

Status and Counters - Accounting Information

Interval(min) : 5
Suppress Empty User : Yes

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

Figure 6-17. 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 6-18. 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 6-19. Example Listing of Active RADIUS Accounting Sessions on the Switch

Changing RADIUS-Server Access Order

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

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

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

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

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

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

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

Figure 6-20. 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.

RADIUS Authentication and Accounting

Changing RADIUS-Server Access Order

```
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 :

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

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

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

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

Figure 6-21. 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 <code>Can't reach RADIUS server < x.x.x.x ></code> , 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.

— This page is intentionally unused —

Configuring RADIUS Server Support for Switch Services

Contents

Overview	7-2
Configuring the RADIUS Server for Per-Port CoS and Rate-Limiting Services	7-3
Viewing the Currently Active Per-Port CoS and Rate-Limiting Configuration Specified by a RADIUS Server	7-4
Configuring and Using RADIUS-Assigned Access Control Lists ...	7-8
Introduction	7-8
Terminology	7-8
Overview of RADIUS-Assigned, Dynamic Port ACLs	7-11
Contrasting Dynamic and Static ACLs	7-13
How a RADIUS Server Applies a Dynamic Port ACL to a Switch Port .. 7-15	
General ACL Features, Planning, and Configuration	7-16
The Packet-filtering Process	7-16
Operating Rules for Dynamic Port ACLs	7-17
Configuring an ACL in a RADIUS Server	7-18
Configuring ACE Syntax in RADIUS Servers	7-21
Configuration Notes	7-22
Configuring the Switch To Support Dynamic Port ACLs	7-24
Displaying the Current Dynamic Port ACL Activity on the Switch	7-25
Event Log Messages	7-28
Causes of Client Deauthentication Immediately After Authenticating	7-29
Monitoring Shared Resources	7-29

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
- Rate-Limiting
- ACLS

Optional Network Management Applications. Per-port CoS and rate-limiting assignments through a RADIUS server are also supported in the ProCurve Manager (PCM) application. Per-port 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 the chapter 6, “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.)

Configuring the RADIUS Server for Per-Port CoS and Rate-Limiting Services

This section provides general guidelines for configuring a RADIUS server to dynamically apply CoS (Class of Service) and Rate-Limiting 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 and Rate-Limiting values have been imposed by a RADIUS server, the CoS and Rate-Limiting 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 7-7.)

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.

Service	Control Method and Operating Notes:
Rate-Limiting on inbound traffic This feature assigns a bandwidth limit 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: 46 (integer = HP) Setting: HP-RATE-LIMIT = <i>< bandwidth-in-Kbps ></i> Note: The CLI command for configuring a rate-limit on a port uses a percentage value. However, using a VSA on a RADIUS server to specify a rate-limit requires the actual Kbps to which you want to limit inbound traffic volume. Thus, to limit in-bound traffic on a gigabit port to 50% of the port's bandwidth capacity requires a VSA setting of 500000 (1,000,000 x 0.5). 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 Rate-Limiting, refer to "Rate-Limiting" in the "Port Traffic Controls" chapter of the <i>Management and Configuration Guide</i> for your switch.

Viewing the Currently Active Per-Port CoS and Rate-Limiting 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. For example, if the switch configuration allows port B1 a rate-limit of 80% of the port's available bandwidth, but the RADIUS server specifies a rate-limit of 50% for a given authenticated client, then the switch shows the RADIUS-imposed rate-limit for that port as long as the authenticated client session is active.

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

These commands display the CoS and Rate-Limiting 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 and Rate-Limiting configuration.

show rate-limit all 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 Rate-Limiting 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	% Curr. Rate Limit Inbound	RADIUS ACL Applied?
B7	Open	1	No-override	No-override	
B8	Closed	1	No-override	No-override	
B9	Open	7		80	
B10	Closed	1	No-override	No-override	

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) and maximum bandwidth for inbound traffic on port B7.

Open indicates that there is an authenticated client session running on port B9. The numeric values (**7** and **80**) are the most recent RADIUS-imposed settings for the CoS (802.1p priority) and maximum bandwidth allowed for inbound traffic on port B9. Refer to the **Note** on page 7-7.

Figure 7-1. Example of Displaying Inbound CoS and Rate-Limiting Imposed by a RADIUS Session

Configuring RADIUS Server Support for Switch Services
 Configuring the RADIUS Server for Per-Port CoS and Rate-Limiting Services

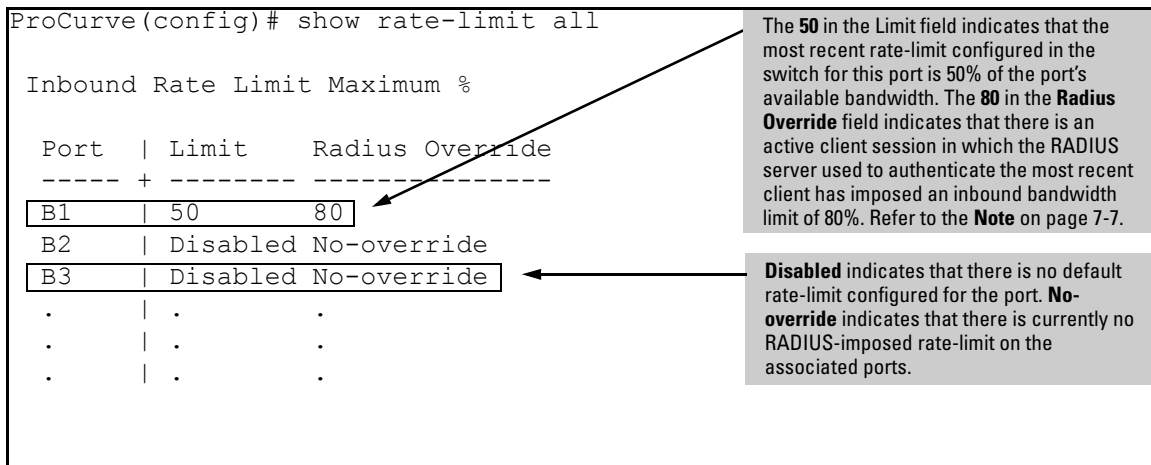


Figure 7-2. Example of Displaying Inbound Rate-Limiting Imposed by a RADIUS Session

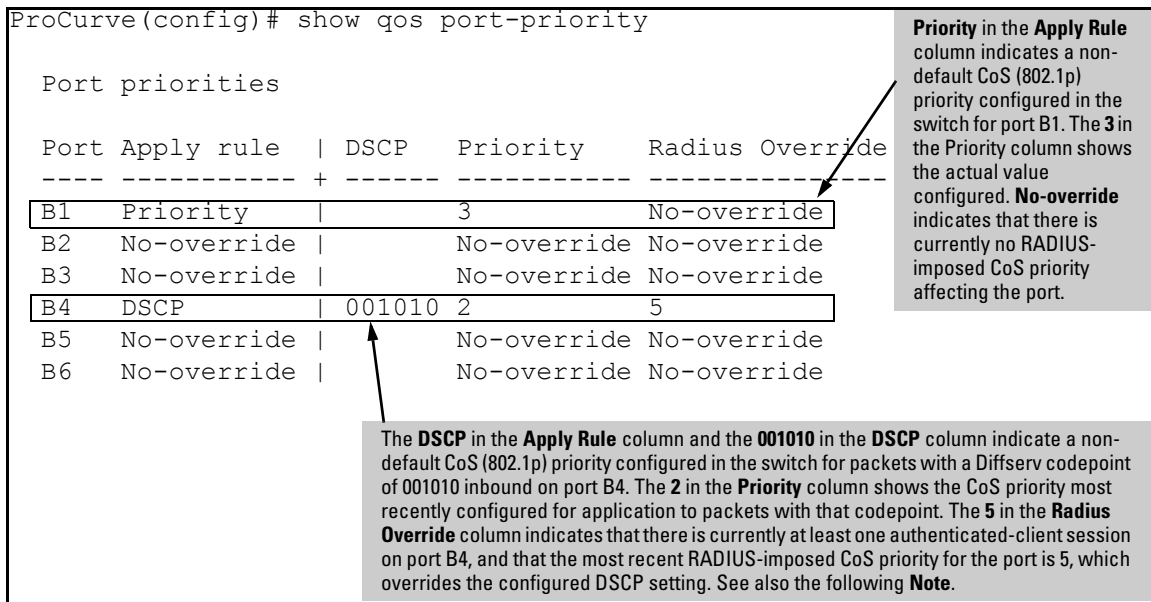


Figure 7-3. Example of Displaying Inbound CoS (802.1p) Priority Imposed by a RADIUS Session

Note

Where multiple clients are currently authenticated on a given port where inbound CoS and Rate-Limiting values have been imposed by a RADIUS server, the port operates with the inbound CoS priority and rate-limit assigned by RADIUS for the most recently authenticated client. Any earlier CoS or rate-limit 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 a rate-limit of 75%, and client “Y” later becomes authenticated with a CoS of 3 and a rate-limit of 50% while the session for client “X” is still active, then the port will operate with a CoS of 3 and a rate-limit of 50% 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 7-2.)

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 the chapter 10, “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:

- **RACL:** an ACL assigned to filter routed traffic entering or leaving the switch on a VLAN. (Separate assignments are required for inbound and outbound traffic.)
- **VACL:** an ACL assigned to filter inbound traffic on a specific VLAN configured on the switch
- **Static Port ACL:** an ACL assigned to filter inbound traffic on a specific switch port
- **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 an RACL, VACL, or 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. (Other, statically configured ACL application types are described in the chapter titled "Access Control Lists (ACLs)" in the *Advanced Traffic Management Guide* for your switch.

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, but VACLs do not filter outbound switched traffic.

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”.

Routed ACL (RACL): An ACL applied to routed traffic that is entering or leaving the switch on a given VLAN. See also “Access Control List”.

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.

VLAN ACL (VACL): An ACL applied to traffic entering the switch on a given VLAN interface. See also “Access Control List”.

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 7-9.

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. Also, dynamic port ACLs can be assigned regardless of whether other ACLs affecting the same port are statically configured on the switch. (For information on statically configured ACLs and application methods, refer to chapter 10, “Access Control Lists (ACLs)”.)

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. Also, applying dynamic port ACLs to ports on the network edge is likely to be less complex than configuring static port and VLAN-based ACLs in the network core to filter unwanted IP traffic that could have been filtered at the edge.

Note

A dynamic port ACL can be applied to a port regardless of whether IP traffic on the port is already being filtered by a static port ACL and/or any VLAN-based ACLs configured on the switch. For more information, refer to “Multiple ACLs on an Interface” on page 10-20.

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 7-2.)

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 7-1, below, highlights several key differences between the static ACLs configurable on switch VLANs and ports, and the dynamic port ACLs that can be assigned to individual ports by a RADIUS server.

Table 7-1. Contrasting Dynamic and Static ACLs

Dynamic Port ACLs	Static Port and VLAN ACLs
Configured in client accounts on a RADIUS server.	Configured on switch ports and VLANs.
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">• selected routed IP traffic (RACLs)• switched or routed IP traffic entering the switch from multiple sources or from unauthenticated sources• IP 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 or VLAN.
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 one each of the following: <ul style="list-style-type: none">• inbound RACL• outbound RACL• VACL• static port ACL
Supports only extended ACLs. (Refer to Terminology.)	Supports standard, extended, and connection-rate ACLs. (Refer to "Configuring and Applying Connection-Rate ACLs" on page 3-19.)

Dynamic Port ACLs	Static Port and VLAN ACLs
<p>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.)</p> <p>Requires client authentication by a RADIUS server configured to dynamically assign an ACL to the client port, based on client credentials.</p> <p>ACEs allow a counter (cnt) option that causes a counter to increment when there is a packet match.</p>	<p>An RACL applied to inbound traffic on a VLAN filters all routed IP traffic entering the switch through a port on that VLAN, as well as any inbound traffic having a DA on the switch itself. An RACL applied to outbound traffic on a VLAN filters all routed IP traffic leaving the switch through a port on that VLAN (and includes routed traffic generated by the switch itself).</p> <p>A VACL applied on a VLAN filters all IP traffic entering the switch through a port on that VLAN.</p> <p>A static port ACL applied on a port filters all traffic entering the switch through that port.</p> <p>No client authentication requirement.</p> <p>ACEs allow a log option that generates a log message whenever there is a packet match with a “deny” ACE.</p>

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. For more information, refer to "Multiple ACLs on an Interface" on page 10-20.

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 10, “Access Control Lists (ACLs)”:

- “Features Common to All ACL Applications” on page 10-22
- “General Steps for Planning and Configuring ACLs” on page 10-24
- “Planning an ACL Application” on page 10-30

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 the chapter 10, “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 and/or belongs to a VLAN for which there is an inbound, VLAN-based ACL configured on the switch, 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. (If the packet is also subject to ACL mirroring, the mirroring action occurs regardless of whether a permit or deny match occurs with any other ACL.)

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 authentication 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 7-18.
- **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 Other, Statically Configured ACLs:** Suppose that port “X” belongs to VLAN “Y” and has a dynamic port ACL assignment from a RADIUS server to filter inbound traffic from an authenticated client. Port “X” is also configured with a static port ACL, and VLAN “Y” is statically configured with a VACL. Any IP traffic entering the switch on port “X” from the client and having a match with a **deny** ACE configured in *any* of these ACLs will be dropped. If an inbound RACL

was also configured on VLAN “Y”, then a **deny** match in the RACL would apply as well to any inbound, routed traffic from the client (and to any inbound, switched traffic having a destination on the switch itself). (If an outbound RACL was also configured on VLAN “Y”, then any outbound, routed IP traffic leaving the switch through the subject port would be filtered by the outbound RACL as well.)

- **Effect of Dynamic Port ACLs on Inbound Traffic for Multiple Clients on the Same Port:** On a port configured for 802.1X *user-based* access where multiple clients are connected, if a given client's authentication results in a dynamic port ACL assignment, then the authentication of any 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 multiple 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 any 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. For more on this topic, refer to “Static Port ACL and Dynamic Port ACL Applications” on page 10-19, and “Multiple ACLs on an Interface” on page 10-20.

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:

```
VENDOR      HP      11
BEGIN-VENDOR HP
ATTRIBUTE   HP-IP-FILTER-RAW 61 STRING
END-VENDOR  HP
```

ProCurve (HP) Vendor-Specific ID

ProCurve (HP) Vendor-Specific Attribute for Dynamic Port ACLs

Note that if you were also using the RADIUS server to administer 802.1p (CoS) priority and/or Rate-Limiting, you would also insert the ATTRIBUTE entries for these functions above the END-VENDOR entry.

Figure 7-4. 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:

```
client 10.10.10.125
nastype = other
secret = 1234
```

Note: The **key** configured in the switch and the **secret** configured in the RADIUS server supporting the switch must be identical. Refer to the chapter titled “RADIUS Authentication and Accounting” in the *Access Security Guide* for your switch.

Figure 7-5. 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 conservation 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 figure 7-6 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”.

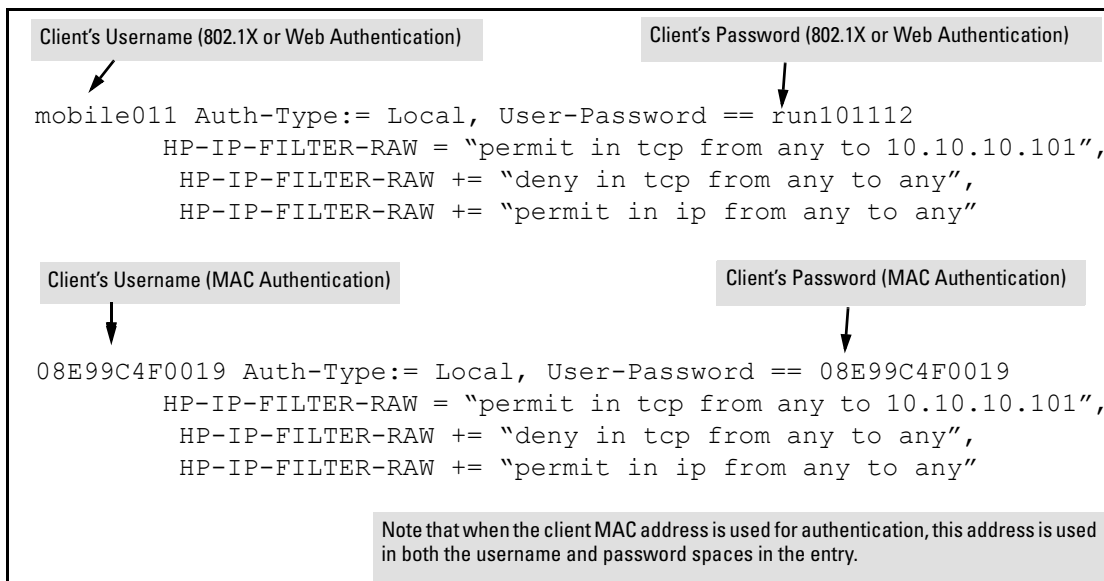


Figure 7-6. Example of Configuring the FreeRADIUS Server To Support ACLs for the Indicated Clients

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 7-22.)

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 6, “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 13, “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 4, “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 4, “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 there are no ACLs currently assigned to any port in < port-list >, executing this command returns only the system prompt. 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.

Configuring RADIUS Server Support for Switch Services

Configuring and Using RADIUS-Assigned Access Control Lists

```
ProCurveSwitch# show access-list radius b1
Radius-configured Port-based ACL for
[Port B1, Client -- 001185C6547D]
[deny in tcp from any to 15.30.248.184 23 cnt]
  Packet Hit Counter : 0
deny in tcp from any to 15.30.248.184 80 cnt
  [Packet Hit Counter : 0]
permit in tcp from any to 15.30.248.184 7
[permit in udp from any to 15.30.248.184 7]
deny in tcp from any to 15.30.248.184 161 cnt
  Packet Hit Counter : 0
deny in udp from any to 15.30.248.184 161 cnt
  Packet Hit Counter : 0
permit in ip from any to any
```

Indicates MAC address identity of the authenticated client on the specified port. This data identifies the client to which the ACL applies.

Lists "deny" ACE for Inbound Telnet (23 = TCP port number) traffic, with counter configured to show the number of matches detected.

Lists current counter for the preceding "Deny" ACE.

Lists "permit" ACEs for inbound TCP and UDP traffic, with no counters configured.

Note that the implicit "deny any/any" included automatically at the end of every ACL is not visible in ACL listings generate by the switch.

Figure 7-7. 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.

% Curr.Rate Limit Inbound: Indicates the status of the current rate-limit setting for inbound traffic.

No-override: No RADIUS-assigned rate-limit is currently active on the indicated port. (For more on rate-limiting, refer to the chapter titled “Port Traffic Controls” in the Management and Configuration Guide for your switch.)

0 - 100: Indicates that the displayed rate-limit 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# show port-access authenticator b1

Port Access Authenticator Status

Port-access authenticator activated [No] : Yes

Port Status   Current   Current   % Curr. Rate   RADIUS ACL
              VLAN ID   Port COS   Limit Inbound  Applied?
-----
B1   Open    1         7              No-override    Yes
B2   Closed  1         No-override    No-override    No
B3   Open    1         No-override    80              Yes
    
```

Figure 7-8. 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.

Message	Meaning
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.
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.
 - The rule limit of 3048 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

Configuring RADIUS Server Support for Switch Services
Configuring and Using RADIUS-Assigned Access Control Lists

subscribed, new RADIUS-based sessions using dynamic port ACLs cannot be authenticated until the necessary resources are released from other applications. For information on determining the current resource availability and usage, refer to the appendix titled “Monitoring Resources” in the *Management and Configuration Guide* for your switch.

Configuring Secure Shell (SSH)

Contents

Overview	8-2
Terminology	8-3
Prerequisite for Using SSH	8-5
Public Key Formats	8-5
Steps for Configuring and Using SSH for Switch and Client Authentication	8-6
General Operating Rules and Notes	8-8
Configuring the Switch for SSH Operation	8-9
1. Assigning a Local Login (Operator) and Enable (Manager) Password	8-9
2. Generating the Switch's Public and Private Key Pair	8-10
3. Providing the Switch's Public Key to Clients	8-12
4. Enabling SSH on the Switch and Anticipating SSH Client Contact Behavior	8-15
5. Configuring the Switch for SSH Authentication	8-18
6. Use an SSH Client To Access the Switch	8-21
Further Information on SSH Client Public-Key Authentication .	8-22
Messages Related to SSH Operation	8-27

Overview

Feature	Default	Menu	CLI	Web
Generating a public/private key pair on the switch	No	n/a	page 8-10	n/a
Using the switch's public key	n/a	n/a	page 8-12	n/a
Enabling SSH	Disabled	n/a	page 8-15	n/a
Enabling client public-key authentication	Disabled	n/a	pages 8-19, 8-22	n/a
Enabling user authentication	Disabled	n/a	page 8-18	n/a

The switches covered in this guide use Secure Shell version 2 (SSHv2) to provide remote access to management functions on the switches via encrypted paths between the switch and management station clients capable of SSH operation.

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

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

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

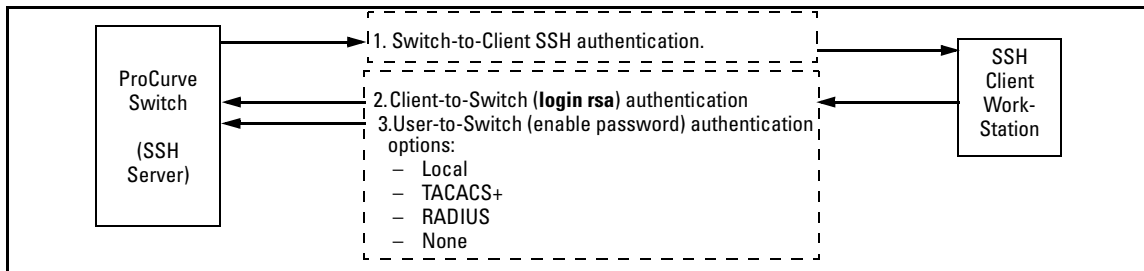


Figure 8-1. Client Public Key Authentication Model

Note

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

Switch SSH and User Password Authentication . This option is a subset of the client public-key authentication shown in figure 8-1. 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 8-1, 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.

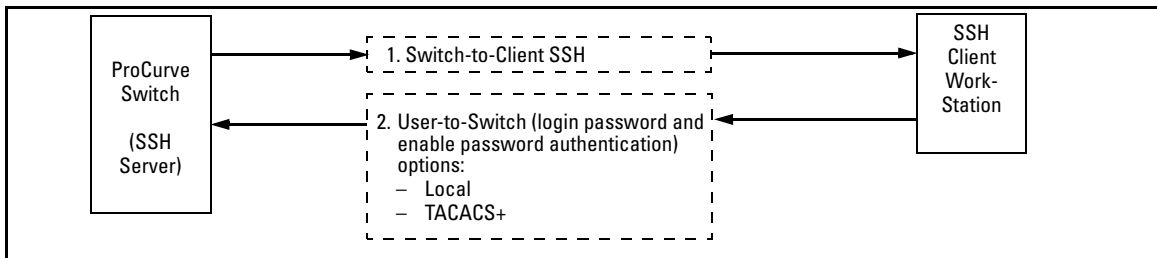


Figure 8-2. Switch/User Authentication

On the switches covered in this guide, SSH supports these data encryption methods:

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

Note

ProCurve switches use RSA keys for internally generated keys. The switch supports both RSA and DSA/DSS keys for clients. All references to either a public or private key mean keys generated using these algorithms, unless otherwise noted

Terminology

- **SSH Server:** An 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 held internally in the switch or by a client.

- **PEM (Privacy Enhanced Mode):** Refers to an ASCII-formatted client public-key that has been encoded for portability and efficiency. SSHv2 client public-keys are typically stored in the PEM format. See figure 8-3 for an example of PEM-encoded ASCII keys.
- **Private Key:** An internally generated key used in the authentication process. A private key generated by the switch is not accessible for viewing or copying. A private key generated by an SSH client application is typically stored in a file on the client device and, together with its public key counterpart, can be copied and stored on multiple devices.
- **Public Key:** An internally generated counterpart to a private key. A device's public key is used to authenticate the device to other devices.
- **Enable Level:** Manager privileges on the switch.
- **Login Level:** Operator privileges on the switch.
- **Local password or username:** A Manager-level or Operator-level password configured in the switch.
- **SSH Enabled:** (1) A public/private key pair has been generated on the switch (**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. Generating the Switch's Public and Private Key Pair” on page 8-10 and “4. Enabling SSH on the Switch and Anticipating SSH Client Contact Behavior” on page 8-15.)

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 8-2), 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 to export public keys. The switch can accept keys in the PEM-Encoded ASCII Format or in the Non-Encoded ASCII format.

```
"Pub Key Gen 21 Dec 2001 12:01"A1B3Hz1y2+orEhYL . . . Q8D8qDH1ozu1c="*** End of Pub Key ***"
```

Comment describing public

Beginning of actual SSHv2 public key in PEM-Encoded

Figure 8-3. Example of Public Key in PEM-Encoded ASCII Format Common for SSHv2 Clients

Steps for Configuring and Using SSH for Switch and Client Authentication

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

Table 8-1. SSH Options

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

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

The general steps for configuring SSH include:

A. Client Preparation

1. Install an SSH client application on a management station you want to use for access to the switch. (Refer to the documentation provided with your SSH client application.)
2. Optional—If you want the switch to authenticate a client public-key on the client:
 - a. Either generate a public/private key pair on the client computer (if your client application allows) or import a client key pair that you have generated using another SSH application.
 - b. Copy the client public key into an ASCII file on a TFTP server accessible to the switch and download the client public key file to the switch. (The client public key file can hold up to 10 client keys.) This topic is covered under “To Create a Client-Public-Key Text File” on page 8-23.

B. Switch Preparation

1. Assign a login (Operator) and enable (Manager) password on the switch (page 8-9).
2. Generate a public/private key pair on the switch (page 8-10).

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 8-12).
4. Enable SSH on the switch (page 8-15).
5. Configure the primary and secondary authentication methods you want the switch to use. In all cases, the switch will use its host-public-key to authenticate itself when initiating an SSH session with a client.
 - SSH Login (Operator) options:
 - Option A:

Primary: Local, TACACS+, or RADIUS password
Secondary: Local password or none. If the primary option is local, the secondary option must be none.
 - Option B:

Primary: Client public-key authentication (**login public-key** — page 8-22)
Secondary: none

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

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

General Operating Rules and Notes

- Public keys generated on an SSH client must be exportable to the switch. The switch can only store 10 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.
- The switch does not support outbound SSH sessions. Thus, if you Telnet from an SSH-secure switch to another SSH-secure switch, *the session is not secure*.
- ▶ With SSH running, the switch allows one console session and up to five other sessions (SSH and/or Telnet).

Configuring the Switch for SSH Operation

SSH-Related Commands in This Section	Page
show ip ssh	8-17
show crypto client-public-key [<manager operator>] [keylist-str] [< babble fingerprint>]	8-25
show crypto host-public-key [< babble fingerprint >]	8-14
show authentication	8-21
crypto key < generate zeroize > ssh [rsa]	8-11
ip ssh	8-16
port < 1 - 65535 default >	8-17
timeout < 5 - 120 >	8-17
aaa authentication ssh	
login < local tacacs radius public-key >	8-18, 8-20
< local none >	8-18
enable < tacacs radius local >	8-18
< local none >	8-18
copy tftp pub-key-file <tftp server IP> <public key file> [<append manager operator>]	8-25
clear crypto client-public-key [keylist-str]	8-26

1. Assigning a 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 8-4. Example of Configuring Local Passwords

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

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

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

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

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 8-14.)

[`fingerprint`]

Displays fingerprints of the switch’s public key in hexadecimal format. (See “Displaying the Public Key” on page 8-14.)

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 AAAAB3NzaClyc2EAAAABIwAAAHEAnAAApdhq13Jynrs7j4lDUm8ivVm81d2mZU5e+YZWp/T6
QzP2RsDDMZLbAHHIBrxPLjW/bRogpYD0lWuVOhTojEVjqeVuXbwmdDnyOgBc06olePwdrbQ+FZevERiA
JYG3C8NCzCRD/djXeI7FmRps8w==
-----
```

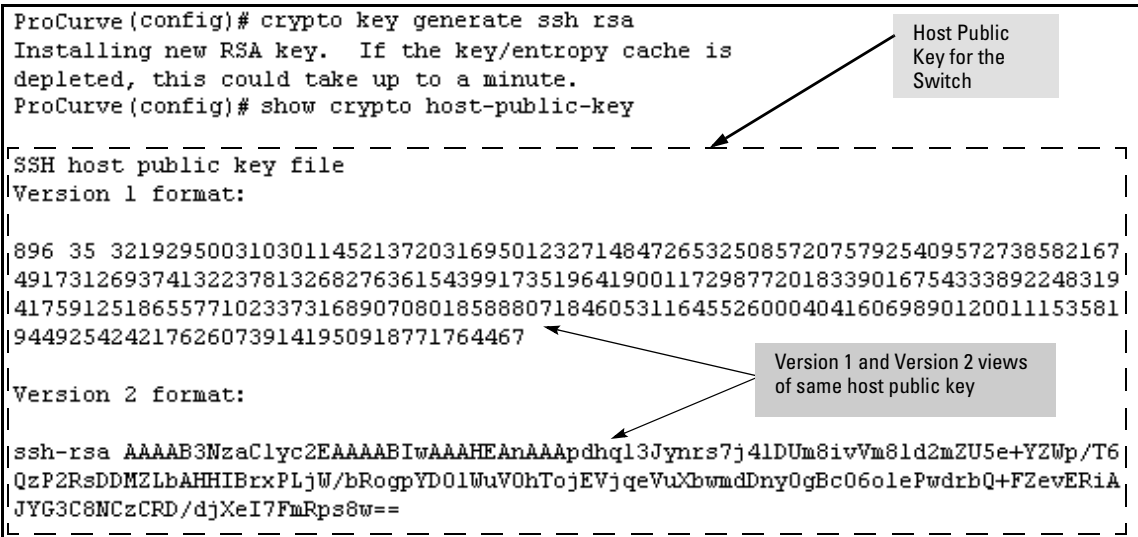


Figure 8-5. Example of Generating a Public/Private Host Key Pair for the Switch

The 'show crypto host-public-key' displays it 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. Providing the Switch's Public Key to Clients

When an SSH client contacts the switch for the first time, the client will challenge the connection unless you have already copied the key into the client's "known host" file. Copying the switch's key in this way reduces the chance that an unauthorized device can pose as the switch to learn your access passwords. The most secure way to acquire the switch's public key for

distribution to clients is to use a direct, serial connection between the switch and a management device (laptop, PC, or UNIX workstation), as described below.

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




Figure 8-6. Example of a Public Key Generated by the Switch

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

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

1. Use a terminal application such as HyperTerminal to display the switch's public key with the **show crypto host-public-key** command (figure 8-5).
2. Bring up the SSH client's "known host" file in a text editor such as Notepad as straight ASCII text, and copy the switch's public key into the file.
3. Ensure that there are no changes or breaks in the text string. (A public key must be an unbroken ASCII string. Line breaks are not allowed. Changes in the line breaks will corrupt the Key.) For example, if you are using Windows® Notepad, ensure that **Word Wrap** (in the **Edit** menu) is disabled, and that the key text appears on a single line.




Figure 8-7. Example of a Correctly Formatted Public Key

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

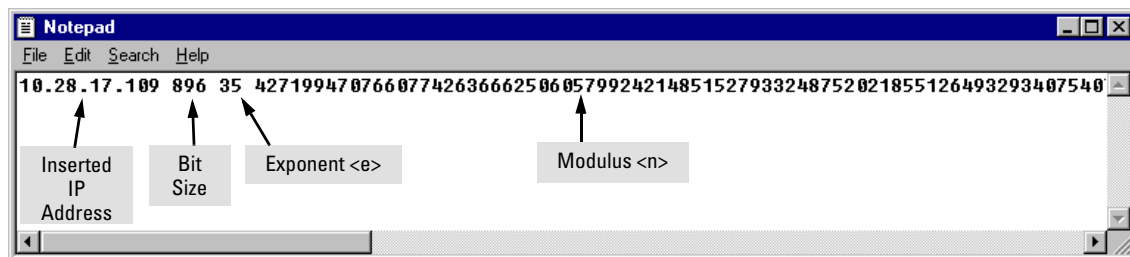


Figure 8-8. Example of a Switch Public Key Edited To Include the Switch's IP Address

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

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

- **Non-encoded ASCII numeric string:** Requires a client ability to display the keys in the "known hosts" file in the ASCII format. This method is tedious and error-prone due to the length of the keys. (See figure 8-7 on page 8-13.)
- **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 8-7 as follows:

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



Figure 8-9. Examples of Visual Phonetic and Hexadecimal Conversions of the Switch's Public Key

The two commands shown in figure 8-9 convert the displayed format of the switch's (host) public key for easier visual comparison of the switch's public key to a copy of the key in a client's "known host" file. The switch has only one RSA host key. The 'babble' and 'fingerprint' options produce two hashes for the key—one that corresponds to the challenge hash you will see if connecting with a v1 client, and the other corresponding to the hash you will see if connecting with a v2 client. These hashes do not correspond to different keys, but differ only because of the way v1 and v2 clients compute the hash of the same RSA key. The switch always uses ASCII version (without babble or fingerprint conversion) of its public key for file storage and default display format.

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

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

Note

Before enabling SSH on the switch you must generate the switch's public/private key pair. If you have not already done so, refer to "2. Generating the Switch's Public and Private Key Pair" on page 8-10.

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. Configuring the Switch for SSH Authentication" on page 8-18.

SSH Client Contact Behavior. At the first contact between the switch and an SSH client, if the switch's public key has not been copied into the client, then the client's first connection to the switch will question the connection and, for security reasons, provide the option of accepting or refusing. If it is safe to assume that an unauthorized device is not using the switch's IP address in an attempt to gain access to the client's data or network, the connection can be accepted. (As a more secure alternative, the client can be directly connected to the switch's serial port to download the switch's public key into the client. See the following Note.)

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. This possibility can be removed by directly connecting the management station to the switch's serial port, using a **show** command to display the switch's public key, and copying the key from the display into a file. This requires a knowledge of where the client stores public keys, plus the knowledge of what key editing and file format might be required by the client application. However, if the first contact attempt between a client and the switch does not pose a security problem, this is unnecessary.

To enable SSH on the switch.

1. Generate a public/private key pair if you have not already done so. (Refer to "2. Generating the Switch's Public and Private Key Pair" on page 8-10.)
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 8-11).

Syntax: [no] ip ssh

Enables or disables SSH on the switch.

[port < 1-65535 | default >]

*The TCP port number for SSH connections (default: 22). **Important:** See "Note on Port Number" on page 8-17.*

[timeout < 5 - 120 >]

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

The **ip ssh key-size** command affects only a per-session, internal server key the switch creates, uses, and discards. This key is not accessible from the user interface. The switch's public (host) key is a separate, accessible key that is always 2048 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 switch 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 12.255.255.255:1873
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 **Note**, below.

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

Figure 8-10. Example of Enabling IP SSH and Listing the SSH Configuration and Status

Caution

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

SSH does not protect the switch from unauthorized access via the web interface, Telnet, SNMP, or the serial port. While web and Telnet access can be restricted by the use of passwords local to the switch, if you are unsure of the security this provides, you may want to disable web-based and/or Telnet access (**no web-management** and **no telnet**). If you need to increase SNMP security, you should use SNMP version 3 only. If you need to increase the security of your web interface see the section on 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. Configuring the Switch for SSH Authentication

Note that all methods in this section result in authentication of the switch's public key by an SSH client. However, only Option B, below results in the switch also authenticating the client's public key. Also, for a more detailed discussion of the topics in this section, refer to "Further Information on SSH Client Public-Key Authentication" on page 8-22

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. Assigning a Local Login (Operator) and Enable (Manager) Password" on page 8-9.

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**. If the primary method is **local**, the secondary method must be **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**. If the primary method is **local**, the secondary method must be **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 8-22.)

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**).*

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**.*

*If the primary access method is **local**, you can only specify **none** for a secondary access method.*

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:

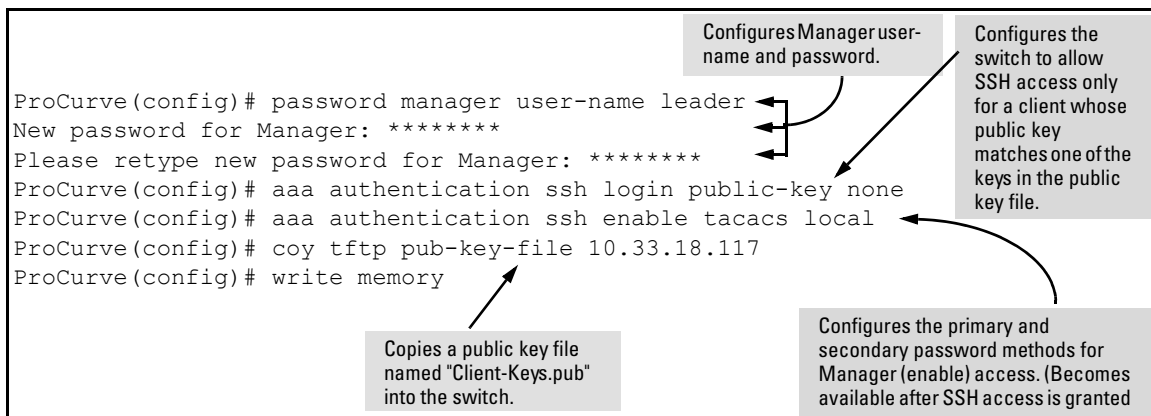


Figure 8-11. Configuring for SSH Access Requiring a Client Public-Key Match and Manager Passwords

Figure 8-12 shows how to check the results of the above commands.

```

ProCurve (config)# show authentication

Status and Counters - Authentication Information

Login Attempts : 3

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

```

Client Key Index Number

```

ProCurve (config)# show crypto client-public-key
0,"Maden name [1024-bit rsa, Local_crypto@Localcrypto, Thu Nov 07 2002 21:25:42]" ssh-rsa AAAAB3NzaClyc2EAAAADAQABAAQgQCz9oNfQxMHUFEC6frSv1Sa4UhlEFznFhQqmgP29HXp6NR/1QOUmACtrFU+QD11EtM/YM9FrN/XvZH/kIXtdEc5exFX/S10tcFafYzI9UjK80dBmQvBGKB
LzVEbCVwlqdAqbkaEX3d/WaPS2xArLCFHsTZhnCvQTZDOGAB1frlcw==
1,"[768-bit rsa, Local_crypto@Localcrypto, Mon Dec 16 2002 23:01:51]" ssh-rsa
AAAAB3NzaClyc2EAAAADAQABAAQgQD0tmzA32JBgeuFJN0iXI3bfooPKZ09JKCPQcXEVk7N+eKf9MOX
vnmfFuEpw/fpqhlvsE66n8FDu7W/B2tKH/tgQLFqx7GiVcxNGhLiNO/pg5AuEym8Enc1Gu/LgAM9daM=

```

Lists the current SSH authentication configuration.

Shows the contents of the public key file downloaded with the `copy tftp` command in figure 8-11. In this example, the file contains two client public-keys.

Figure 8-12. SSH Configuration and Client-Public-Key Listing From Figure 8-11

6. Use an SSH Client To Access the Switch

Test the SSH configuration on the switch to ensure that you have achieved the level of SSH operation you want for the switch. If you have problems, refer to "RADIUS-Related Problems" in the Troubleshooting chapter of the *Management and Configuration Guide* for your switch.

Further Information on SSH Client Public-Key Authentication

The section titled “5. Configuring the Switch for SSH Authentication” on page 8-18 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. Configuring the Switch for SSH Authentication” on page 8-18.)

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 from step 6 and compares it to the client's hash version. If they match, then the client is authenticated. Otherwise, the client is denied access.

Using client public-key authentication requires these steps:

1. Generate a public/private key pair for each client you want to have SSH access to the switch. This can be a separate key for each client or the same key copied to several clients.
2. Copy the public key for each client into a client-public-key text file.
3. Use **copy tftp** to copy the client-public-key file 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
↓	↙	↓	↘
<pre>1024 35 11407406661701446907963803652840180539127043745111482882509285550110168603082 6038959146896306569035982041222025542543282764329943344032963504381021098947647460564 5572227682031607648603664020534703408371002884293231503492265409355321119922465153140 745413543765609589968291386053556814705585051025488575846923smith@support.cairns.com</pre>			

Figure 8-13. Example of a Client Public Key

Notes

Comments in public key files, such as **smith@support.cairns.com** in figure 8-13, may appear in a SSH client application's generated public key. While such comments may help to distinguish one key from another, they do not pose any restriction on the use of a key by multiple clients and/or users.

Public key illustrations such as the key shown in figure 8-13 usually include line breaks as a method for showing the whole key. However, in practice, line breaks in a public key will cause errors resulting in authentication failure.

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

Property	Supported Value	Comments
Key Format	ASCII	See figure 8-7 on page 8-13. 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 individually on a TFTP server to which the switch has access. Terminate all client public-keys in the file except the last one with a <CR><LF>.

Note on Public Keys

The actual content of a public key entry in a public key file is determined by the SSH client application generating the key. (Although you can manually add or edit any comments the client application adds to the end of the key, such as the **smith@support.cairns.com** at the end of the key in figure 8-13 on page 8-23.)

Syntax: copy tftp pub-key-file <ip-address> <filename> [append | manager | operator>]

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

*The **append** option adds the key(s) for operator access.*

*The **manager** option replaces the key(s) for manager access; follow with the 'append' option to add the key(s).*

*The **operator** option replaces the key(s) for operator access (default); follow with the 'append' option to add the key(s).*

show crypto client-public-key [<manager | operator>] [keylist-str] [babble | fingerprint]

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

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

*The **fingerprint** option converts the key data to hexadecimal hashes that are for the same purpose.*

*The **keylist-str** selects keys to display (comma-delimited list).*

*The **manager** option allows you to select manager public keys*

*The **operator** option allows you to select operator public keys.*

Note

Beginning with software release K_12_XX or later, **copy usb pub-key file** can also be used as a method for copying a public key file to the switch.

Configuring Secure Shell (SSH)

Further Information on SSH Client Public-Key Authentication

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
IyVEbCVwlqdAqbkaEX3d/WaPS2xArLCFHsTZhnCvQTZDOGAB1frlcw==
1."[768-bit rsa, Jamie_wilson@Jamiewilson, Mon Dec 16 2002 23:01:51]" ssh-rsa
AAAAB3NzaClyc2EAAAADAQABAAQgYQD0tmzA32JBgeuFJN0iXI3bfooPKZ09JKCPQcXEVk7N+eKf9MOX
vnmfFuBpw/fpqhlvsE66n8FDu7W/B2tKH/tqQLFcx7GiVcxNGhLiN0/pq5AuEym8Enc1Gu/LgAM9daM=
```

Key Index Number

Figure 8-14. Example of Copying and Displaying a Client Public-Key File Containing Two Different Client Public Keys for the Same Client

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

- You can remove the existing client public-key file or specific keys by executing the **clear crypto public-key** command.

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

Syntax: clear crypto public-key 3
Deletes the entry with an index of 3 from the client-public-key file on the switch.

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

- If an SSH client's public key matches the switch's client-public-key file, allow that client access to the switch. If there is not a public-key match, then deny access to that client.

Syntax: aaa authentication ssh login public-key none

Allows SSH client access only if the switch detects a match between the client's public key and an entry in the client-public-key file most recently copied into the switch.

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.	File transfer did not occur. 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.	File transfer did not occur. 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 8-17.
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.

Configuring Secure Shell (SSH) Messages Related to SSH Operation

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

Configuring Secure Socket Layer (SSL)

Contents

Overview	9-2
Terminology	9-3
Prerequisite for Using SSL	9-5
Steps for Configuring and Using SSL for Switch and Client Authentication	9-5
General Operating Rules and Notes	9-6
Configuring the Switch for SSL Operation	9-7
1. Assigning a Local Login (Operator) and Enable (Manager) Password	9-7
2. Generating the Switch's Server Host Certificate	9-9
To Generate or Erase the Switch's Server Certificate with the CLI	9-10
Comments on certificate fields.	9-11
Generate a Self-Signed Host Certificate with the Web browser interface	9-13
Generate a CA-Signed server host certificate with the Web browser interface	9-15
3. Enabling SSL on the Switch and Anticipating SSL Browser Contact Behavior	9-17
Using the CLI interface to enable SSL	9-19
Using the web browser interface to enable SSL	9-19
Common Errors in SSL setup	9-21

Overview

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

The switches covered in this guide use Secure Socket Layer Version 3 (SSLv3) and support for Transport Layer Security(TLSv1) to provide remote web access to the switches via encrypted paths between the switch and management station clients capable of SSL/TLS operation.

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 the switches covered in this guide is based on the OpenSSL software toolkit. For more information on OpenSSL, visit www.openssl.com.

Server Certificate authentication with User Password

Authentication . This option is a subset of full certificate authentication of the user and host. It occurs only if the switch has SSL enabled. As in figure 9-1, 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 9-1. Switch/User Authentication

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

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

Note:

ProCurve Switches use RSA public key algorithms and Diffie-Hellman, and all references to a key mean keys generated using these algorithms unless otherwise noted

Terminology

- **SSL Server:** An 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 browser documentation for additional details

B. Switch Preparation

1. Assign a login (Operator) and enable (Manager) password on the switch. (page 9-7)
2. Generate a host certificate on the switch. (page 9-9)
 - 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 9-17)
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

Configuring the Switch for SSL Operation

SSL-Related CLI Commands in This Section	Page
web-management ssl	page 9-19
show config	page 9-19
show crypto host-cert	page 9-12
crypto key	
generate cert [rsa] <512 768 1024>	page 9-10
zeroize cert	page 9-10
crypto host-cert	
generate self-signed [arg-list]	page 9-10
zeroize	page 9-10

1. Assigning a 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 ProCurve Web Browser Interface” in the *Management and Configuration Guide* for your switch.

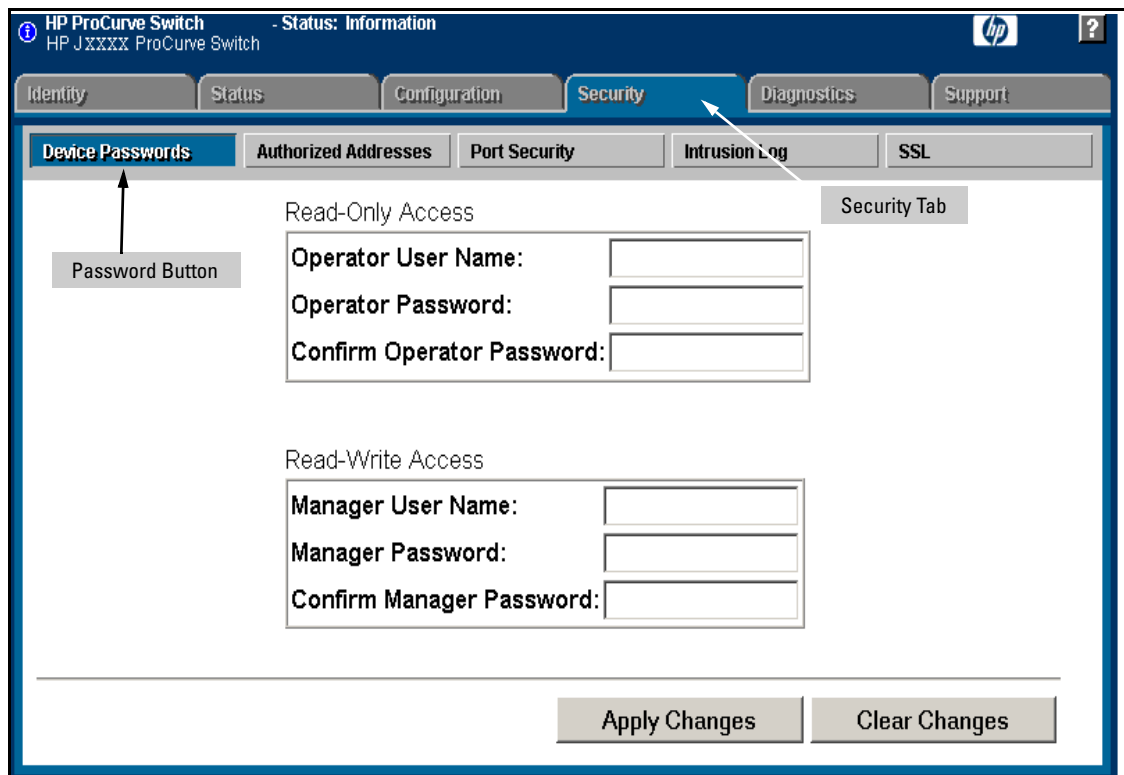


Figure 9-2. Example of Configuring Local Passwords

1. Proceed to the security tab and select device passwords button.
2. Click in the appropriate box in the Device Passwords window and enter user names and passwords. You will be required to repeat the password strings in the confirmation boxes.

Both the user names and passwords can be up to 16 printable ASCII characters.
3. Click on **[Apply Changes]** button to activate the user names and passwords.

2. Generating the Switch's Server Host Certificate

You must generate a server certificate on the switch before enabling SSL. The switch uses this server certificate, along with a dynamically generated session key pair to negotiate an encryption method and session with a browser trying to connect via SSL to the switch. (The session key pair mentioned above is not visible on the switch. It is a temporary, internally generated pair used for a particular switch/client session, and then discarded.)

The server certificate is stored in the switch's flash memory. The server certificate should be added to your certificate folder on the SSL clients who you want to have access to the switch. Most browser applications automatically add the switch's host certificate to their certificate folder on the first use. This method does allow for a security breach on the first access to the switch. (Refer to the documentation for your browser application.)

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

Note:

There is usually a fee associated with receiving a verified certificate and the valid dates are limited by the root certificate authority issuing the certificate.

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

Removing (zeroizing) the switch's certificate key pair or certificate render the switch unable to engage in SSL operation and automatically disables SSL on the switch. (To verify whether SSL is enabled, execute **show config**.)

To Generate or Erase the Switch's Server Certificate with the CLI

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

CLI commands used to generate a Server Host Certificate.

Syntax: `crypto key generate cert [rsa] < 512 | 768 | 1024 >`

Generates a key pair for use in the certificate.

`crypto key zeroize cert`

Erases the switch's certificate key and disables SSL operation.

`crypto host-cert generate self-signed [arg-list]`

Generates a self signed host certificate for the switch. If a switch certificate already exists, replaces it with a new certificate. (See the Note, above.)

`crypto host-cert zeroize`

Erases the switch's host certificate and disables SSL operation.

To generate a host certificate from the CLI:

- i. Generate a certificate key pair. This is done with the **crypto key generate cert** command. The default key size is 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 9-1, “Certificate Field Descriptions” describes these arguments.

Table 9-1. 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]: 01/01/2002
Validity end date   [01/01/2003]: 01/01/2004
Common name        [10.255.255.255]: 10.255.255.255
Organization        [Company Name]: Hewlett Packard
Organizational unit [Dept Name]: ProCurve Network
City or location    [City]: Roseville
State name          [State]: Ca
Country code        [US]: US
  
```

The diagram shows three text boxes with arrows pointing to the CLI output:

- Generate New Key**: Points to the command `crypto key generate cert 512`.
- Generate New Certificate**: Points to the command `crypto host-cert generate self-signed`.
- Enter certificate Arguments**: Points to the interactive prompts for validity dates, common name, organization, organizational unit, city, state, and country code.

Figure 9-3. Example of Generating a Self-Signed Server Host certificate on the CLI for the Switch.

Notes

“Zeroizing” the switch’s server host certificate or key automatically disables SSL (sets **web-management ssl** to **No**). Thus, if you zeroize the server host certificate or key and then generate a new key and server certificate, you must also re-enable SSL with the **web-management ssl** command before the switch can resume SSL operation.

CLI Command to view host certificates.

Syntax: show crypto host-cert

Displays switch’s host certificate

To view the current host certificate from the CLI you use the **show crypto host-cert** command.

For example, to display the new server host certificate:

```
ProCurve(config)#show crypto host-cert ← Show host certificate command
Version: 1 (0x0)
Serial Number: 0 (0x0)
Issuer: CN=10.255.255.255, L=Roseville, ST=Ca, C=US, O=Hewlett Packard, OU=ProCurve Network
Validity
  Not Before: Jan 1 00:00:00 2002 GMT
  Not After : Jan 1 23:59:59 2004 GMT
Subject: CN=10.255.255.255, L=Roseville, ST=Ca, C=US, O=Hewlett Packard, OU=ProCurve Network
Subject Public Key Info:
  Public Key Algorithm: rsaEncryption
  RSA Public Key: (512 bit)
  Modulus (512 bit):
    00:db:18:4b:ce:3e:7d:5a:90:d8:a5:50:d5:2a:e9:
    60:78:d1:35:82:e9:27:71:5d:45:8d:0a:b9:b4:55:
    65:c7:d1:1c:4e:30:5e:20:a6:2d:62:9c:4c:cd:40:
    a0:6a:0b:cb:1c:ce:90:1c:2c:ad:26:fc:0b:07:ae:
    db:11:65:d6:47
  Exponent: 35 (0x23)
Signature Algorithm: md5WithRSAEncryption
d6:d0:98:6b:b9:a5:54:96:d9:be:fa:b9:99:f9:d8:6f:94:42:
30:ea:c4:1d:88:e6:7b:19:18:22:84:f6:8c:ea:46:d7:ab:42:
26:48:77:0c:60:57:8c:33:bc:08:d8:f7:c6:1f:ef:15:b7:24:
f3:fa:92:b1:1f:7d:9e:c1:fd:83

MD5 Fingerprint: C969 E196 49C3 4609 AFC6 BDE1 2087 00A7
SHA1 Fingerprint: 93C7 0753 F805 26DC 4E39 EAF2 9C18 174F 7A63 E3C5
```

Figure 9-4. 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 ProCurve 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. Proceed to the Security tab then the SSL button. The SSL configuration screen is split up into two halves. The left half is used in creating a new certificate key pair and (self-signed / CA-signed) certificate. The right half displays information on the currently installed certificate.
- ii. Select the Generate Certificate button.
- iii. Select Self signed certificate in the type box.
- iv. Select the RSA key size desired. If you do not wish to generate a new key then just select current from the list.
- v. Fill in remaining certificate arguments (refer to “To Generate or Erase the Switch’s Server Certificate with the CLI” on page 9-10).
- vi. Click on the **[Apply Changes]** button to generate a 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:




Figure 9-5. Self-Signed Certificate generation via SSL Web Browser Interface Screen

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

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

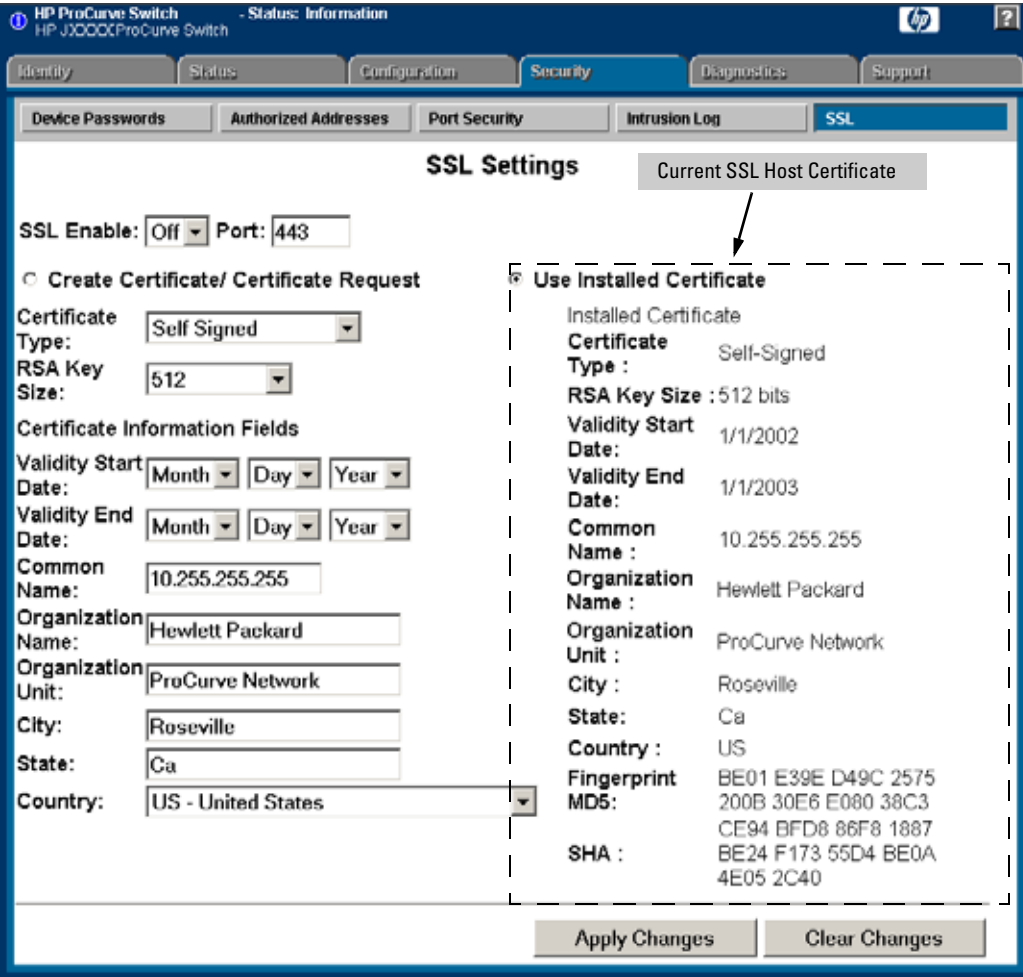


Figure 9-6. Web browser Interface showing current SSL Host Certificate

Generate a CA-Signed server host certificate with the Web browser interface

To install a CA-Signed server host certificate from the web browser interface. For more information on how to access the web browser interface, refer to the chapter titled “Using the ProCurve 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 select the **[SSL]** button
- ii. Select the **Create Certificate/Certificate Request** radio button.
- iii. Select **Create CA Request** from the **Certificate Type** drop-down list.
- iv. Select the key size from the RSA Key Size drop-down list. If you wish to re-use the current certificate key, select **Current** from the **RSA Key Size** drop-down list.
- v. Fill in remaining certificate arguments (Refer to “Comments on certificate fields.” on page 9-11.)
- 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 non- PEM encoded certificate request reply.
- vii. After the certificate authority processes your request and sends you a certificate reply (that is, an installable certificate), copy and paste it into the lower text box.
- viii. Click on the **[Apply Changes]** button to install the certificate.




Figure 9-7. Request for Verified Host Certificate Web Browser Interface Screen

3. Enabling SSL on the Switch and Anticipating SSL Browser Contact Behavior

The **web-management ssl** command enables SSL on the switch and modifies parameters the switch uses for transactions with clients. After you enable SSL, the switch can authenticate itself to SSL enabled browsers. If you want to disable SSL on the switch, use the **no web-management ssl** command.

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

When configured for SSL, the switch uses its host certificate to authenticate itself to SSL clients, however unless you disable the standard ProCurve 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 9-20.*

show config

*Shows status of the SSL server. When enabled **web-management ssl** will be present in the config list.*

To enable SSL on the switch

1. Generate a Host certificate if you have not already done so. (Refer to “2. Generating the Switch’s Server Host Certificate” on page 9-9.)
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 9-10).

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 9-8. 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 switches are 49, 80, 1506, and 1513.

Caution

SSL does not protect the switch from unauthorized access via the Telnet, SNMP, or the serial port. While Telnet access can be restricted by the use of passwords local to the switch, if you are unsure of the security this provides, you may want to disable Telnet access (**no telnet**). If you need to increase SNMP security, use SNMP version 3 only for SNMP access. Another security measure is to use the Authorized IP Managers feature described in the switch's *Security Guide*. To protect against unauthorized access to the serial port (and the Clear button, which removes local password protection), keep physical access to the switch restricted to authorized personnel.

Common Errors in SSL setup

Error During	Possible Cause
Generating host certificate on CLI	You have not generated a certificate key. (Refer to "CLI commands used to generate a Server Host Certificate" on page 9-10.)
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 9-13.)</p> <p>You may be using a reserved TCP port. (Refer to "Note on Port Number" on page 9-20.)</p>
Unable to Connect with SSL	<p>You may not have SSL enabled (Refer to "3. Enabling SSL on the Switch and Anticipating SSL Browser Contact Behavior" on page 9-17.)</p> <p>Your browser may not support SSLv3 or TLSv1 or it may be disabled. (Refer to the documentation provided for your browser.)</p>

— This page is intentionally unused —

Access Control Lists (ACLs)

Contents

Introduction	10-4
Overview of Options for Applying ACLs on the Switch	10-5
Static ACLS	10-5
Dynamic Port ACLs	10-5
Terminology	10-10
Overview	10-15
Types of IP ACLs	10-15
ACL Applications	10-15
RACL Applications	10-16
VACL Applications	10-18
Static Port ACL and Dynamic Port ACL Applications	10-19
Multiple ACLs on an Interface	10-20
Features Common to All ACL Applications	10-22
General Steps for Planning and Configuring ACLs	10-24
ACL Operation	10-26
Introduction	10-26
The Packet-filtering Process	10-27
Planning an ACL Application	10-30
IP Traffic Management and Improved Network Performance	10-30
Security	10-32
Guidelines for Planning the Structure of an ACL	10-32
ACL Configuration and Operating Rules	10-33
How an ACE Uses a Mask To Screen Packets for Matches	10-36
What Is the Difference Between Network (or Subnet) Masks and the Masks Used with ACLs?	10-36
Rules for Defining a Match Between a Packet and an Access Control Entry (ACE)	10-37

Configuring and Assigning an ACL	10-41
Overview	10-41
General Steps for Implementing ACLs	10-41
Options for Permit/Deny Policies	10-42
ACL Configuration Structure	10-42
Standard ACL Structure	10-43
Extended ACL Configuration Structure	10-45
ACL Configuration Factors	10-46
The Sequence of Entries in an ACL Is Significant	10-46
Allowing for the Implied Deny Function	10-48
A Configured ACL Has No Effect Until You Apply It to an Interface	10-48
You Can Assign an ACL Name or Number to an Interface Even if the ACL Does Not Exist in the Switch's Configuration	10-48
Using the CLI To Create an ACL	10-49
General ACE Rules	10-49
Using CIDR Notation To Enter the ACL Mask	10-50
Configuring Standard ACLs	10-51
Configuring Named, Standard ACLs	10-53
Creating Numbered, Standard ACLs	10-56
Configuring Extended ACLs	10-60
Configuring Named, Extended ACLs	10-62
Configuring Numbered, Extended ACLs	10-74
Adding or Removing an ACL Assignment On an Interface	10-81
Filtering Routed IP Traffic	10-81
Filtering IP Traffic Inbound on a VLAN	10-82
Filtering Inbound IP Traffic Per Port	10-84
Deleting an ACL	10-85
Editing an Existing ACL	10-86
Using the CLI To Edit ACLs	10-86
General Editing Rules	10-86
Sequence Numbering in ACLs	10-87
Inserting an ACE in an Existing ACL	10-88
Deleting an ACE from an Existing ACL	10-90
Resequencing the ACEs in an ACL	10-91

Attaching a Remark to an ACE	10-92
Operating Notes for Remarks	10-95
Displaying ACL Configuration Data	10-96
Display an ACL Summary	10-97
Display the Content of All ACLs on the Switch	10-98
Display the RACL and VACL Assignments for a VLAN	10-99
Display Static Port ACL Assignments	10-100
Displaying the Content of a Specific ACL	10-101
Display All ACLs and Their Assignments in the Routing Switch Startup-Config File and Running-Config File	10-103
Creating or Editing ACLs Offline	10-104
Creating or Editing an ACL Offline	10-104
The Offline Process	10-104
Example of Using the Offline Process	10-105
Enable ACL “Deny” Logging	10-109
Requirements for Using ACL Logging	10-109
ACL Logging Operation	10-110
Enabling ACL Logging on the Switch	10-111
General ACL Operating Notes	10-113

Introduction

An Access Control List (ACL) is a list of one or more Access Control Entries (ACEs) specifying the criteria the switch uses to either permit (forward) or deny (drop) IP packets traversing the switch's interfaces. This chapter describes how to configure, apply, and edit ACLs in a network populated with the switches covered by this guide, and how to monitor ACL actions.

Feature	Default	CLI
Standard ACLs	None	10-51
Extended ACLs	None	10-60
Enable or Disable an ACL	n/a	10-81
Display ACL Data	n/a	10-96
Delete an ACL	n/a	10-85
Configure an ACL from a TFTP Server	n/a	10-104
Enable ACL Logging	n/a	10-111

IP filtering with ACLs can help improve network performance and restrict network use by creating policies for:

- **Switch Management Access:** Permits or denies in-band management access. This includes limiting and/or preventing the use of designated protocols that ride on top of IP, such as TCP, UDP, IGMP, ICMP, and others. Also included are the use of precedence and ToS criteria, and control for application transactions based on source and destination IP addresses and transport layer port numbers.
- **Application Access Security:** Eliminates unwanted IP traffic in a path by filtering IP packets where they enter or leave the switch on specific VLAN interfaces.

ACLs can filter IP traffic to or from a host, a group of hosts, or entire subnets.

Notes

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

ACLs on the switches covered by this manual do not screen non-IP traffic such as AppleTalk and IPX.

Overview of Options for Applying ACLs on the Switch

To apply ACL filtering, assign a configured ACL to the interface on which you want the IP traffic filtering to occur. VLAN and routed IP traffic ACLs can be applied statically using the switch configuration. Port traffic ACLs can be applied either statically or dynamically (using a RADIUS server).

Static ACLS

Static ACLs are configured on the switch. To apply a static ACL, you must assign it to an interface (VLAN or port). The switch supports three static ACL applications:

Routed IP Traffic ACL (RACL). An RACL is an ACL configured on a VLAN to filter routed IP traffic entering or leaving the switch on that interface, as well as IP traffic having a destination on the switch itself. (Except for filtering IP traffic to an IP address on the switch itself, RACLs can operate only while IP routing is enabled. Refer to “Notes on IP Routing” on page 10-25.)

VLAN ACL (VACL). A VACL is an ACL configured on a VLAN to filter IP traffic entering the switch on that VLAN interface and having a destination on the same VLAN.

Static Port ACL. A static port ACL is an ACL configured on a port to filter IP traffic entering the switch on that port, regardless of whether the IP traffic is routed, switched, or addressed to a destination on the switch itself.

Dynamic Port ACLs

A dynamic port ACL is configured on a RADIUS server for assignment to a given port when the server authenticates a specific client on that port. When the client is authenticated, the ACL configured for that client on the server is assigned to the port and applied to the IP traffic received inbound on that port from the authenticated client. When the client session ends, the ACL is removed from the port. The switch allows as many dynamic port ACLs on a port as it allows authenticated clients.

Access Control Lists (ACLs)

Overview of Options for Applying ACLs on the Switch

Note

This chapter describes the ACL applications you can statically configure on the switch. For information on dynamic port ACLs assigned by a RADIUS server, refer to the chapter 7, “Configuring RADIUS Server Support for Switch Services”.

Table 10-1. Command Summary for Standard ACLs

Action	Command(s)	Page
Create a Standard, Named ACL <i>or</i> Add an ACE to the End of an Existing Standard, Named ACL	ProCurve(config)# ip access-list standard < name-str > ProCurve(config-std-nacl)# < deny permit > < any host <SA > SA/< mask-length > SA < mask >> ¹ [log] ²	10-53
Create a Standard, Numbered ACL <i>or</i> Add an ACE to the End of an Existing Standard, Numbered ACL	ProCurve(config)# access-list < 1-99 > < deny permit > < any host <SA > SA/< mask-length > SA < mask >> [log] ²	10-56
Use a Sequence Number To Insert an ACE in a Standard ACL	ProCurve(config)# ip access-list standard < name-str 1-99 > ProCurve(config-std-nacl)# 1-2147483647 < deny permit > < any host <SA > SA/< mask-length > SA < mask >> ¹ [log] ²	10-87
Use an ACE's Sequence Number To Delete the ACE from a Standard ACL	ProCurve(config)# ip access-list standard < name-str 1-99 > ProCurve(config-std-nacl)# no < 1-2147483647 >	10-90
Resequence the ACEs in a Standard ACL	ProCurve(config)# ip access-list resequence < name-str 1-99 > < 1-2147483647 > < 1-2147483646 >	10-91
Enter or Remove a Remark from a Standard ACL	ProCurve(config)# ip access-list standard < name-str 1-99 > ProCurve(config-ext-nacl)# [remark < remark-str > no < 1-2147483647 > remark]	10-92 10-94

*For numbered, standard ACLs only, the following **remark** commands can be substituted for the above:*

```
ProCurve(config)# access-list < 1 - 99 > remark < remark-str >  
ProCurve(config)# [no] access-list < 1 - 99 > remark
```

Delete a Standard ACL	ProCurve(config)# no ip access-list standard < name-str 1-99 >	10-85
-----------------------	--	-------

For numbered, standard ACLs, the following command can be substituted for the above:

ProCurve(config)# access-list < 1 - 99 > remark < remark-str >

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

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

Access Control Lists (ACLs)

Overview of Options for Applying ACLs on the Switch

Table 10-2. Command Summary for Extended ACLs

Action	Command(s)	Page
Create an Extended, Named ACL	ProCurve(config)# ip access-list extended < name-str 100-199 > ProCurve(config-std-nacl)# < deny permit >	10-62
<i>or</i>	< ip ip-protocol ip-protocol-nbr >	
Add an ACE to the End of an Existing, Extended ACL	< any host <SA > SSA/< mask-length > SA < mask >> ¹ < any host < DA > DA/< mask-length > DA < mask >> ¹ < tcp udp > < any host <SA > SA/< mask-length > SA < mask >> ¹ [comparison-operator < value >] < any host <DA > DA/< mask-length > DA < mask >> ¹ [comparison-operator < value >] [established] < igmp > < any host <SA > SA/< mask-length > SA < mask >> ¹ < any host < DA > DA/< mask-length > DA < mask >> ¹ [igmp-packet-type] < icmp > < any host <SA > SA/< mask-length > SA < mask >> ¹ < any host < DA > DA/< mask-length > DA < mask >> ¹ [[< 0 - 255 > [0 - 255]] icmp-message] [precedence < priority >] [tos < tos- setting >] [log] ²	
Create an Extended, Numbered ACL	ProCurve(config)# access-list < 100-199 > < deny permit > < ip-options tcp/udp-options igmp-options icmp-options > [precedence < priority >] [tos < tos- setting >] [log] ²	10-74
<i>or</i>		
Add an ACE to the End of an Existing, Numbered ACL	Note: Uses the same IP, TCP/UDP, IGMP, and ICMP options as shown above for "Create an Extended, Named ACL".	
Insert an ACE by Assigning a Sequence Number	ProCurve(config)# ip access-list extended < name-str 100-199 > ProCurve(config-ext-nacl)# 1-2147483647 < deny permit > Uses the options shown above for "Create an Extended, Named ACL".	10-88
Delete an ACE by Specifying Its Sequence Number	ProCurve(config)# ip access-list extended < name-str 100-199 > ProCurve(config-std-nacl)# no < 1-2147483647 >	10-90
Resequence the ACEs in an ACL	ProCurve(config)# ip access-list resequence < name-str 100-199 > < 1-2147483647 > < 1-2147483646 >	10-91

¹The mask can be in either dotted-decimal notation (such as 0.0.15.255) or CIDR notation (such as /20).
²The [log] function applies only to "deny" ACLs, and generates a message only when there is a "deny" match.

Action	Command(s)	Page
Enter or Remove a Remark	ProCurve(config)# ip access-list extended < name-str 100-199 >	10-92
	ProCurve(config-ext-nacl)# [remark < remark-str > no remark]	10-94
	<i>For numbered, extended ACLs only, the following remark commands can be substituted for the above:</i>	
	ProCurve(config)# access-list < 100 - 199 > remark < remark-str >	
	ProCurve(config)# [no] access-list < 100 - 199 > remark	
Delete an Extended ACL	ProCurve(config)# no ip access-list extended < name-str 100-199 >	10-85
	<i>For numbered, extended ACLs only, the following command can also be used:</i>	
	ProCurve(config)# no access-list < 100 - 199 >	

Table 10-3. Command Summary for Enabling, Disabling, and Displaying ACLs

Enable or Disable an RACL	ProCurve(config)# [no] vlan < vid > ip access-group < identifier > < in out >	10-81
Enable or Disable a VACL	ProCurve(config)# [no] vlan < vid > ip access-group < identifier > < vlan >	
Enable or Disable a Static Port ACL	ProCurve(config)# [no] interface < port-list Trkx > access-group < identifier > in ProCurve(eth-< port-list > Trkx >)# [no] ip access-group < identifier > in	
Displaying ACL Data	ProCurve(config)# show access-list ProCurve(config)# show access-list < acl-identifier > ProCurve(config)# show access-list config ProCurve(config)# show access-list vlan < vid > ProCurve(config)# show access-list radius	10-96

Terminology

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

- source IP address and mask (standard and extended ACLs)
- destination IP address and mask (extended ACLs only)
- either of the following:
 - all IP traffic
 - IP traffic of a specific IP protocol (extended ACLs only)
(In the cases of TCP, UDP, ICMP, and IGMP, the criteria can include either all IP traffic of the protocol type or only the IP traffic of a specific sub-type within the protocol.)
- option to log packet matches with **deny** ACEs
- optional use of IP precedence and ToS settings (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” ACE. ACLs can be used to filter IP traffic and to select IP traffic to be monitored (mirrored). ACL types include “standard” and “extended”. See “Standard ACL” and “Extended ACL”. For filtering IP traffic, both can be applied in any of the following ways:

- **RACL:** an ACL assigned to filter routed IP traffic entering or leaving the switch on a VLAN. (Separate assignments are required for inbound and outbound IP traffic.)
- **VACL:** an ACL assigned to filter inbound IP traffic on a specific VLAN configured on the switch
- **Static Port ACL:** an ACL assigned to filter inbound IP traffic on a specific switch port
- **Dynamic Port ACL:** dynamic ACL assigned to a port by a RADIUS server to filter inbound IP traffic from an authenticated client on that port

An ACL can be configured on a VLAN as an RACL or VACL (or both), and on a port (or static trunk) as a static port ACL. (Dynamic port ACLs are configured on a RADIUS server.)

See also “ACL Mirroring”.

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 an alphanumeric string or a number in the range of 1 to 99. An *extended* ACL ID can have either an alphanumeric string or a number in the range of 100 to 199. See also “Identifier”.

Note: RADIUS-assigned ACLs are identified by client authentication data and do not use the ACL ID strings described here.

ACL Mask: Follows any IP address (source or destination) 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). See also “How an ACE Uses a Mask To Screen Packets for Matches” on page 10-36.)

CIDR: This is the acronym for Classless Inter-Domain Routing.

Connection-Rate ACL: An optional feature used with Connection-Rate filtering based on virus-throttling technology. For more information, refer to the chapter 3, “Virus Throttling”.

DA: The acronym used in text to represent *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 IP packet for which there is a match within an applicable ACL.

Dynamic Port ACL: An ACL assigned by a RADIUS server to a port to filter inbound IP traffic from a client authenticated by the server for that port. A dynamic port ACL filters all inbound IP traffic, regardless of whether it is switched or routed. When the client session ends, the dynamic port ACL for that client is removed from the port.

Extended ACL: This type of Access Control List uses layer-3 IP criteria composed of source and destination IP addresses and (optionally) TCP/UDP port, ICMP, IGMP, precedence, or ToS criteria to determine whether there is a match with an IP packet. Except for RADIUS-assigned ACLs, which use client credentials for identifiers, extended ACLs require an alphanumeric name or an identification number (ID) in the range of 100 - 199.

identifier: The term used in ACL syntax statements to represent either the name or number by which the ACL can be accessed. See also **name-str**.
Note: RADIUS-assigned ACLs are identified by client authentication data and do not use the identifiers described in this chapter.

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

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

- Routed ACL (RACL): Inbound traffic is any IP packet entering the switch on a VLAN interface (or a subnet in a multinetted VLAN) with a destination IP address (DA) that is for any of the following:
 - an external device on a different VLAN or subnet than the interface on which it arrived
 - an IP address configured on the switch itself
 - a broadcast

Note that, except for IP traffic addressed to the switch itself, and outbound IP traffic generated by the switch, routing must be configured on the switch to enable support for RACL applications.

- VLAN ACL (VACL): Inbound traffic is any IP packet entering the switch on a VLAN interface (or a subnet in a multinetted VLAN).
- Static Port ACL: Inbound traffic is any IP packet entering the switch on the port.
- Dynamic Port ACL: Where a RADIUS server has authenticated a client and assigned an ACL to the port to filter the client’s IP traffic, inbound traffic is any IP packet entering the switch from that client.

name-str: The term used in extended ACL syntax statements to represent the “name string”; the alphanumeric string used to identify the ACL. See also **identifier** and **ACL-ID**.

Named ACL: An ACL created with the **ip access-list < extended | standard > < name-str >** command and then populated using the **< deny | permit >** command in the Named ACL (**nacl**) CLI context. (Refer to “Entering the “Named ACL” (nacl) Context” on page 10-53.)

Numbered ACL: An ACL created and initially populated by using the **access-list < 1-99 | 100 - 199 >** command. (Refer to “Creating or Adding to a Standard, Numbered ACL” on page 10-57.) After a numbered ACL has been created, the switch manages it in the same way as a named ACL, meaning that it can be applied and edited in the same way as a named ACL.

Outbound Traffic: For defining the points where the switch applies an RACL to filter IP traffic, outbound traffic is routed IP traffic *leaving the switch* through a VLAN interface (or a subnet in a multinetted VLAN). “Outbound traffic” can also apply to switched IP traffic leaving the switch on a VLAN interface, but VACLs do not filter outbound switched IP traffic. (Refer also to “ACL Applications” on page 10-15.)

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

Permit Any Forwarding: An ACE configured with this action causes the switch to forward all routed IP packets that have not been permitted or denied by earlier ACEs in the list. In a standard ACL, this is **permit any**. In an extended ACL, it is **permit ip any any**.

RACL: See “Routed ACL”.

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

remark-str: The term used in ACL syntax statements to represent the variable “remark string”; a set of alphanumeric characters you can include in a remark in an ACL. A remark string can include up to 100 characters and must be delimited by single or double quotes if any spaces are included in the string.

Routed ACL (RACL): An ACL applied to routed IP traffic that is entering or leaving the switch on a given VLAN. See also “Access Control List”.

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 a standard ACE, this is the IP address used by the ACE to determine whether there is a match between a packet and the ACE. In an extended ACE, this is the first of two IP addresses used by the ACE to determine whether there is a match between a packet and the ACE. See also “DA”.

seq-#: The term used in ACL syntax statements to represent the sequence number variable used to insert an ACE within an existing list. The range allowed for sequence numbers is 1 - 2147483647.

Standard ACL: This type of access control list uses the layer-3 IP criteria of source IP address to determine whether there is a match with an IP packet. Except for RADIUS-assigned ACLs, standard ACLs require an alphanumeric name or an identification number (ID) in the range of 1-99. See also **identifier** on page 10-12.

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

VACL: See “VLAN ACL”.

VLAN ACL (VACL): An ACL applied to all IP traffic entering the switch on a given VLAN interface. See also “Access Control List”.

Wildcard: The part of a mask that indicates the bits in a packet’s IP addressing that do not need to match the corresponding bits specified in an ACL. See also **ACL Mask** on page 10-11.

Overview

Types of IP ACLs

A permit or deny policy for IP traffic you want to filter can be based on source IP address alone, or on source IP address plus other IP factors.

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

Extended ACL: Use an extended ACL when simple IP source address restrictions do not provide the sufficient IP traffic selection criteria needed on an interface. Extended ACLs allow use of the following criteria:

- source and destination IP address combinations
- IP protocol options

Extended, named ACLs also offer an option to permit or deny IP connections using TCP for applications such as Telnet, http, ftp, and others.

Connection-Rate ACL. An optional feature used with Connection-Rate filtering based on virus-throttling technology. Refer to the chapter 3, “Virus Throttling”.

ACL Applications

ACL filtering is applied to IP traffic as follows:

- Routed ACL (RACL)— on a VLAN configured with an RAACL:
 - routed IP traffic entering or leaving the switch. (Routing can be between different VLANs or between different subnets in the same VLAN. IP routing *must* be enabled.)
 - routed IP traffic having a destination address (DA) on the switch itself. In figure 10-1 on page 10-17, this is any of the IP addresses shown in VLANs “A”, “B”, and “C”. (IP routing need not be enabled.)

- outbound traffic generated by the switch itself.
- VLAN ACL (VACL): on a VLAN configured with a VACL, any inbound IP traffic, regardless of whether it is switched or routed. On a multi-netted VLAN, this includes all inbound IP traffic from any subnet.
- Static port ACL: any inbound IP traffic on that port.
- Dynamic port ACL: on a port having an ACL assigned by a RADIUS server to filter an authenticated client's IP traffic, any inbound IP traffic from that client

(For information on RADIUS-assigned ACLs, refer to chapter 7, “Configuring RADIUS Server Support for Switch Services”.)

- ACL Mirroring: applies an ACL to a port or VLAN to mirror selected IP traffic to a mirror destination. In this context, a **permit** ACE means to mirror the specified IP traffic; a **deny** ACE means to avoid mirroring. (A **log** keyword in a **deny** ACE is ignored when the associated ACL is used for mirroring.) Refer to “Local and Remote Traffic Mirroring” in the appendix titled “Monitoring and Analyzing Switch Operation” in the *Management and Configuration Guide* for your switch.
- Connection-Rate ACL: An optional feature used with Connection-Rate filtering based on virus-throttling technology. Refer to the chapter 3, “Virus Throttling”.

RACL Applications

RACLs filter routed IP traffic entering or leaving the switch on VLANs configured with the “in” and/or “out” ACL option

```
vlan < vid > ip access-group < identifier > < in | out >
```

For example, in figure 10-1:

- You would assign either an inbound ACL on VLAN 1 or an outbound ACL on VLAN 2 to filter a packet routed between subnets on different VLANs; that is, from the workstation 10.28.10.5 on VLAN 1 to the server at 10.28.20.99 on VLAN 2. (An outbound ACL on VLAN 1 or an inbound ACL on VLAN 2 would not filter the packet.)
- Where multiple subnets are configured on the same VLAN, then you can use either inbound or outbound ACLs to filter routed IP traffic between the subnets on the VLAN if the traffic source and destination IP addresses are on devices external to the switch.




Figure 10-1. Example of RACL Filter Applications on Routed IP Traffic

Notes

The switch allows one inbound RACL assignment and one outbound RACL assignment configured per VLAN. This is in addition to any other ACL assigned to the VLAN or to any ports on the VLAN. You can use the same RACL or different RACLs to filter inbound and outbound routed IP traffic on a VLAN.

RACLs do not filter IP traffic that remains in the same subnet from source to destination (switched IP traffic) unless the destination IP address (DA) or source IP address (SA) is on the switch itself.

VACL Applications

VACLs filter any IP traffic entering the switch on a VLAN configured with the “VLAN” ACL option.

```
vlan < vid > ip access-group < identifier > vlan
```

For example, in figure 10-2, you would assign a VACL to VLAN 2 to filter all inbound switched or routed IP traffic received from clients on the 10.28.20.0 network. In this instance, routed IP traffic received on VLAN 2 from VLANs 1 or 3 would not be filtered by the VACL on VLAN 2.

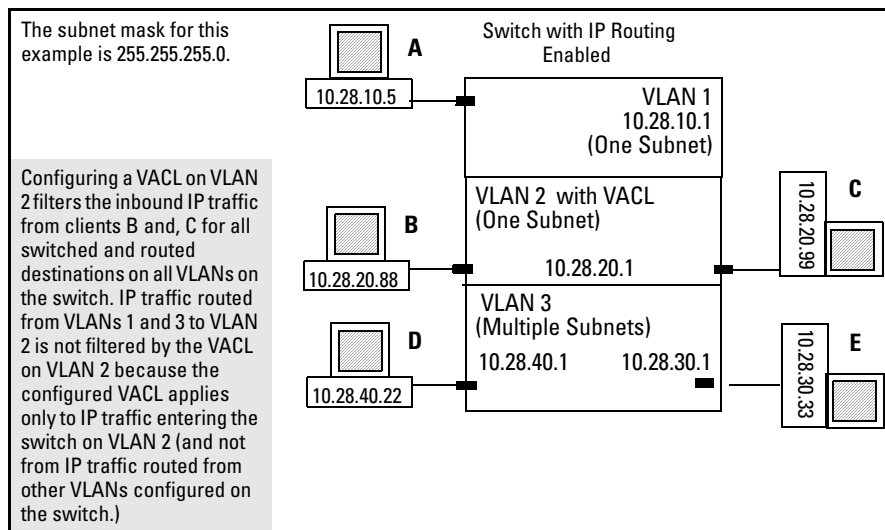


Figure 10-2. Example of VACL Filter Applications on IP Traffic Entering the Switch

Note

The switch allows one VACL assignment configured per VLAN. This is in addition to any other ACL applications assigned to the VLAN or to ports in the VLAN.

Static Port ACL and Dynamic Port ACL Applications

- **Static Port ACL:** filters any IP traffic inbound on the designated port, regardless of whether it is switched or routed.
- **Dynamic (RADIUS-assigned) Port ACL:** filters IP traffic inbound from the client whose authentication resulted in the ACL assignment to the designated port. For example, client “A” connects to a given port and is authenticated by a RADIUS server. Because the server is configured to assign an ACL to the port used by the authenticated client, all IP traffic inbound on the port from client “A” is filtered.

Effect of Dynamic Port ACLs When Multiple Clients Are Using the Same Port. Some network configurations may allow multiple clients to authenticate through a single port where a RADIUS server assigns a separate, dynamic port ACL in response to each client’s authentication on that port. In such cases, a given client’s inbound traffic will be allowed only if the RADIUS authentication response for that client includes a dynamic port ACL. For example, in figure 10-3 (below), clients A through D authenticate through the same port (B1) on the 5400zl switch.




Figure 10-3. Example of Multiple Clients Authenticating Through a Single Port

In this case, the RADIUS server must be configured to assign a dynamic port ACL to port B1 each time any of the clients authenticates on the port.

802.1X User-Based and Port-Based Applications. User-Based 802.1X access control allows up to 32 individually authenticated clients on a given port. However, port-based access control does not set a client limit, and requires only one authenticated client to open a given port (and is recommended for applications where only one client at a time can connect to the port).

- If you configure 802.1X *user-based* security on a port and the RADIUS response includes a dynamic port ACL for at least one authenticated client, then the RADIUS response for all other clients authenticated on the port must also include a dynamic port ACL. Traffic on the port from any client that authenticates without the RADIUS server including a dynamic port ACL in its response will be dropped and the client will be de-authenticated.
- Using 802.1X *port-based* security on a port where the RADIUS response includes a dynamic port ACL, only the first client to authenticate can use the port. Traffic from other clients will be dropped.

Multiple ACLs on an Interface

Multiple ACL Assignments Allowed. The switch allows multiple ACL applications on an interface (subject to internal resource availability). This means that a port belonging to a given VLAN “X” can simultaneously be subject to all of the following:

- One VACL for any IP traffic for VLAN “X” entering the switch through the port.
- One static port ACL for any IP traffic entering the switch on the port.
- One dynamic (RADIUS-assigned) port ACL applied to inbound IP traffic for each authenticated client on the port
- One connection-rate ACL for inbound IP traffic for VLAN “X” on the port (if the port is configured for connection-rate filtering). (Refer to chapter 3, “Virus Throttling”.)
- ACL mirroring per VLAN, port, and trunk interface (Refer to “Local and Remote Traffic Mirroring” in the appendix titled “Monitoring and Analyzing Switch Operation” in the *Management and Configuration Guide* for your switch.)

- One inbound and one outbound RACL filtering routed IP traffic moving through the port for VLAN “X”. (Also applies to inbound, switched traffic on VLAN “X” that has a destination on the switch itself.”

Note

In cases where an RACL and any type of port or VLAN ACL are filtering traffic entering the switch, the *switched* traffic explicitly permitted by the port or VLAN ACL is not filtered by the RACL (except when the traffic has a destination on the switch itself). However, *routed* traffic explicitly permitted by the port or VLAN ACL (and any switched traffic having a destination on the switch itself) must also be explicitly permitted by the RACL, or it will be dropped.

Also, a switched packet is not affected by an outbound RACL assigned to the VLAN on which the packet exits from the switch.

A Packet Must Have a Match with a “Permit” ACE in All Applicable ACLs Assigned to an Interface. On a given interface where multiple ACLs apply to the same traffic, a packet having a match with a **deny** ACE in any applicable ACL on the interface (including an implicit **deny any**) will be dropped.

For example, suppose the following is true:

- Port A10 belongs to VLAN 100.
- A static port ACL is configured on port A10.
- A VACL is configured on VLAN 100.
- An RACL is also configured for inbound, routed traffic on VLAN 100.

An inbound, *switched* packet entering on port A10, with a destination on port A12, will be screened by the static port ACL and the VACL, regardless of a match with any **permit** or **deny** action. A match with a **deny** action (including an implicit deny) in either ACL will cause the switch to drop the packet. (If the packet has a match with explicit **deny** ACEs in multiple ACLs and the log option is included in these ACEs, then a separate log event will occur for each match.) The switched packet will not be screened by the RACL.

However, suppose that VLAN 2 in figure 10-4 (page 10-22) is configured with the following:

- A VACL permitting IP traffic having a destination on the 10.28.10.0 subnet

- An RACL that denies inbound IP traffic having a destination on the 10.28.10.0 subnet

In this case, no IP traffic received on the switch from clients on the 10.28.20.0 subnet will reach the 10.28.10.0 subnet, even though the VACL allows such traffic. This is because the **deny** in the RACL causes the switch to drop the traffic regardless of whether any other VACLs permit the traffic.

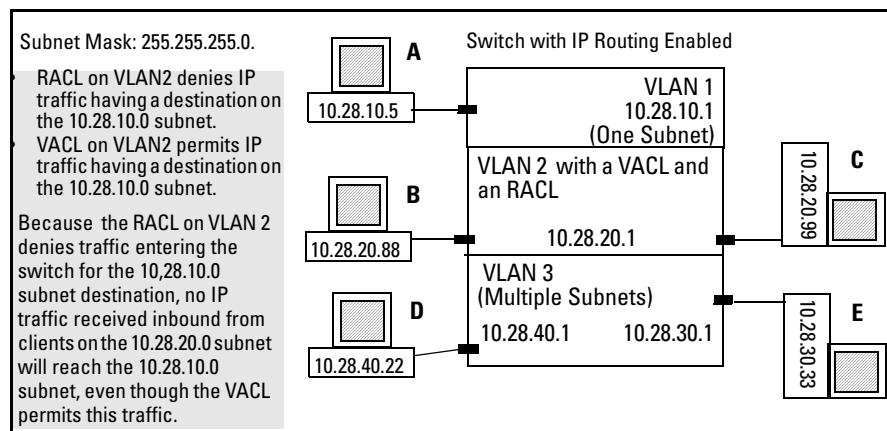


Figure 10-4. Example of Order of Application for Multiple ACLs on an Interface

Exception for Mirrored IP Traffic. If ACL mirroring is configured along with one or more of the above ACL applications on the same interface, the mirroring action occurs regardless of the effect of other ACLs on the packets that match the mirror criteria. This means, for example, that if a dynamic port ACL denies a packet that also meets the mirror ACL criteria for forwarding to the configured mirror destination, the packet will be mirrored even though it will not be forwarded to its intended destination.

Exception for Connection-Rate Filtering. Connection-rate filtering can be configured along with one or more other ACL applications on the same interface. In this case, a connection-rate match for a **filter** action is carried out according to the configured policy, regardless of whether any other ACLs on the interface have a match for a **deny** action. Also, if a connection-rate filter permits (**ignore** action) a packet, it can still be denied by another ACL on the interface.

Features Common to All ACL Applications

- Any ACL can have multiple entries (ACEs).

- You can apply any one ACL to multiple interfaces.
- All ACEs in an ACL configured on the switch are automatically sequenced (numbered). For an existing ACL, entering an ACE without specifying a sequence number automatically places the ACE at the end of the list. Specifying a sequence number inserts the ACE into the list at the correct sequential location.
 - Automatic sequence numbering begins with “10” and increases in increments of 10. You can renumber the ACEs in an ACL and also change the sequence increment between ACEs.
 - The CLI **remark** command option allows you to enter a separate comment for each ACE.
- A source or destination IP address and a mask, together, can define a single host, a range of hosts, or all hosts.
- Every ACL populated with one or more explicit ACEs includes an Implicit Deny as the last entry in the list. The switch applies this action to any packets that do not match other criteria in the ACL. (For standard ACLs, the Implicit Deny is **deny any**. For extended ACLs, it is **deny ip any any**.)
- In any ACL, you can apply an ACL log function to ACEs that have an explicit “deny” action. The logging occurs when there is a match on a “deny” ACE (except when the ACL is used for mirroring). The switch sends ACL logging output to Syslog, if configured, and, optionally, to a console session.

You can create ACLs for the switch configuration using either the CLI or a text editor. The text-editor method is recommended when you plan to create or modify an ACL that has more entries than you can easily enter or edit using the CLI alone. Refer to “Creating or Editing ACLs Offline” on page 10-104.

General Steps for Planning and Configuring ACLs

1. Identify the ACL application to apply. As part of this step, determine the best points at which to apply specific ACL controls. For example, you can improve network performance by filtering unwanted IP traffic at the edge of the network instead of in the core. Also, on the switch itself, you can improve performance by filtering unwanted IP traffic where it is inbound to the switch instead of outbound.

IP Traffic Source	ACL Application
IP traffic from a specific, authenticated client	dynamic port ACL (RADIUS-assigned ACL) for inbound IP traffic from an authenticated client on a port*
IP traffic entering the switch on a specific port	static port ACL (static-port assigned) for any inbound IP traffic on a port from any source
switched or routed IP traffic entering the switch on a specific VLAN	VACL (VLAN ACL)
routed IP traffic entering or leaving the switch on a specific VLAN	RACL (routed ACL)

*For more on this option, refer to chapter 7, "Configuring RADIUS Server Support for Switch Services", and also to the documentation for your RADIUS server.)

2. Identify the IP traffic types to filter:
 - The SA and/or the DA of IP traffic you want to permit or deny. This can be a single host, a group of hosts, a subnet, or all hosts.
 - Any IP traffic of a specific protocol type (0-255)
 - Any TCP traffic (only) for a specific TCP port or range of ports, including optional control of connection traffic based on whether the initial request should be allowed
 - Any UDP traffic (only) or UDP traffic for a specific UDP port
 - Any ICMP traffic (only) or ICMP traffic of a specific type and code
 - Any IGMP traffic (only) or IGMP traffic of a specific type
 - Any of the above with specific precedence and/or ToS settings
3. Design the ACLs for the control points (interfaces) you have selected. Where you are using explicit "deny" ACEs, you can optionally use the ACL logging feature for notification that the switch is denying unwanted packets.
4. Configure the ACLs on the selected switches.

5. Assign the ACLs to the interfaces you want to filter, using the ACL application (static port ACL, VACL, or RACL) appropriate for each assignment. (For RADIUS-assigned ACLs, refer to the Note in the table in step 1 on page 10-24.)
6. If you are using an RACL, ensure that IP routing is enabled on the switch.
7. Test for desired results.

For more details on ACL planning considerations, refer to “Planning an ACL Application” on page 10-30.

Notes on IP Routing

To activate a RACL to screen inbound IP traffic for routing between subnets, assign the RACL to the statically configured VLAN on which the traffic enters the switch. Also, ensure that IP routing is enabled. Similarly, to activate a RACL to screen routed, outbound IP traffic, assign the RACL to the statically configured VLAN on which the traffic exits from the switch. A RACL configured to screen inbound IP traffic with a destination IP address on the switch itself does not require routing to be enabled. (ACLs do not screen outbound IP traffic generated by the switch, itself.) Refer to “ACL Screening of IP Traffic Generated by the Switch” on page 10-113.)

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 interfaces, and offer these traffic filtering options:

- Any IP traffic inbound on a port.
- Any IP traffic inbound on a VLAN.
- Routed IP traffic entering or leaving the switch on a VLAN. (Note that ACLs do not screen traffic at the internal point where traffic moves between VLANs or subnets within the switch. Refer to “ACL Applications” on page 10-15.)

The following table lists the range of interface options:

Interface	ACL Application	Application Point	Filter Action
Port	Static Port ACL (switch configured)	inbound on the switch port	any inbound IP traffic
	Dynamic Port ACL ¹ (RADIUS assigned)	inbound on the switch port used by authenticated client	any inbound IP traffic from the authenticated client
VLAN	VACL	entering the switch on the VLAN	any inbound IP traffic
		entering the switch on the VLAN	routed IP traffic entering the switch and any IP traffic with a destination on the switch itself
	exitting from the switch on the VLAN	routed IP traffic exiting from the switch	

¹This chapter describes ACLs statically configured on the switch. For information on dynamic port ACLs assigned by a RADIUS server, refer to the chapter 7, “Configuring RADIUS Server Support for Switch Services”.

²Supports one inbound and/or one outbound RACL. When both are used, one RACL can be assigned to filter both inbound and outbound, or different RACLs can be assigned to filter inbound and outbound.

Note

After you assign an ACL to an interface, the default action on the interface is to implicitly deny any IP traffic that is not specifically permitted by the ACL. (This applies only in the direction of traffic flow filtered by the ACL.)

The Packet-filtering Process

Sequential Comparison and Action. When an ACL filters a packet, it sequentially compares each ACE's filtering criteria to the corresponding data in the packet until it finds a match. The action indicated by the matching ACE (deny or permit) is then performed on the packet.

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

Example. Suppose the ACL in figure 10-5 is assigned to filter the IP traffic from an authenticated client on a given port in the switch:

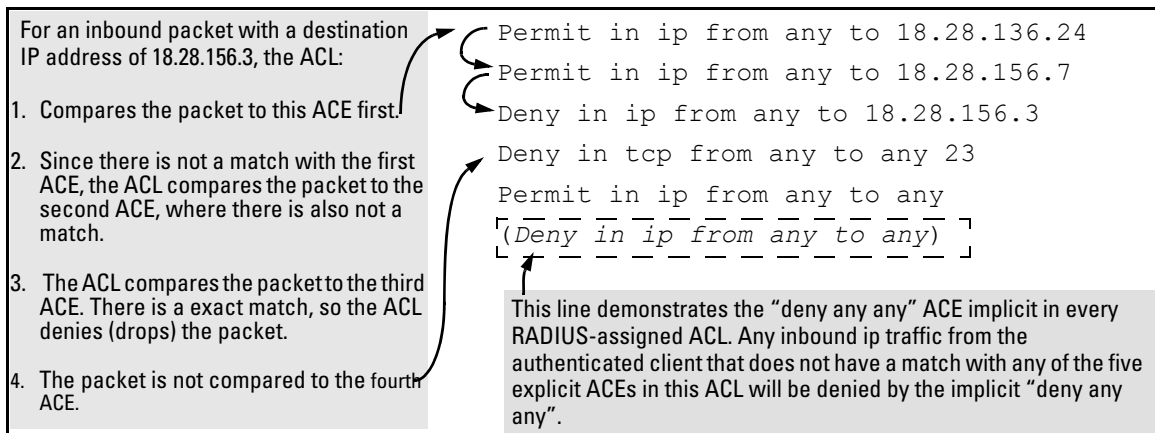


Figure 10-5. Example of Sequential Comparison

As shown above, the ACL tries to apply the first ACE in the list. If there is not a match, it tries the second ACE, and so on. When a match is found, the ACL invokes the configured action for that entry (permit or drop the packet) and

no further comparisons of the packet are made with the remaining ACEs in the list. This means that when an ACE whose criteria matches a packet is found, the action configured for that ACE is invoked, and any remaining ACEs in the ACL are ignored. *Because of this sequential processing, successfully implementing an ACL depends in part on configuring ACEs in the correct order for the overall policy you want the ACL to enforce.*




Figure 10-6. The Packet-Filtering Process in an ACL with *N* Entries (ACEs)

Note

The order in which an ACE occurs in an ACL is significant. For example, if an ACL contains six ACEs, but the first ACE allows Permit Any forwarding, then the ACL permits all IP traffic, and the remaining ACEs in the list do not apply, even if they specify criteria that would make a match with any of the IP traffic permitted by the first ACE.

For example, suppose you want to configure an ACL on the switch (with an ID of "Test-02") to invoke these policies for routed IP traffic entering the switch on VLAN 12 :

1. Permit all inbound IP traffic from IP address 10.11.11.42.
2. Deny *only* the inbound Telnet traffic from address 10.11.11.101.
3. Permit *only* inbound Telnet traffic from IP address 10.11.11.33.
4. Deny *all other* inbound IP traffic.

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

```

ip access-list extended "Test-02"

  1 10 permit ip 10.11.11.42 0.0.0.0 0.0.0.0 255.255.255.255

  2 20 deny tcp 10.11.11.101 0.0.0.0 0.0.0.0 255.255.255.255 eq 23

  3 30 permit ip 10.11.11.101 0.0.0.0 0.0.0.0 255.255.255.255

  4 40 permit tcp 10.11.11.33 0.0.0.0 0.0.0.0 255.255.255.255 eq 23

  5 < Implicit Deny >
exit
ProCurve(config)# vlan 12 ip access-group Test-02 in
  
```

<p>1. Permits IP traffic from source address 10.11.11.42. 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.</p>	<p>4. Permits Telnet traffic from source address 10.11.11.33. 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.</p>
<p>2. Denies Telnet traffic from source address 10.11.11.101. 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.</p>	<p>5. This entry does not appear in an actual ACL, but is implicit as the last entry in every ACL. Any packets that do not match any of the criteria in the ACL's preceding entries will be denied (dropped), and will not cross VLAN 12.</p>
<p>3. Permits any IP traffic from source address 10.11.11.101. 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.</p>	

Figure 10-7. Example of How an ACL Filters Packets

It is important to remember that all ACLs configurable on the switch include an implicit **deny ip any any**. That is, IP packets that the ACL does not *explicitly* permit or deny will be *implicitly* denied, and therefore dropped instead of forwarded on the interface. If you want to preempt the implicit deny so that packets not explicitly denied by other ACEs in the ACL will be permitted, insert an explicit “permit any” as the last ACE in the ACL. Doing so permits any packet not explicitly denied by earlier entries. (Note that this solution does not apply in the preceding example, where the intention is for the switch to forward only explicitly permitted packets routed on VLAN 12.

Planning an ACL Application

Before creating and implementing ACLs, you need to define the policies you want your ACLs to enforce, and understand how the ACL assignments will impact your network users.

Note

All IP traffic entering the switch on a given interface is filtered by all ACLs configured for inbound traffic on that interface. For this reason, an inbound packet will be denied (dropped) if it has a match with either an implicit or explicit **deny** in *any* of the inbound ACLs applied to the interface. (This does not apply to IP traffic leaving the switch because only one type of ACL—an RACL—can be applied, and only to routed IP traffic.)

(Refer to “Multiple ACLs on an Interface” on page 10-20.)

IP Traffic Management and Improved Network Performance

You can use ACLs to block IP traffic from individual hosts, workgroups, or subnets, and to block access to VLANs, subnets, devices, and services. Traffic criteria for ACLs include:

- Switched and/or routed IP traffic
- Any IP traffic of a specific protocol type (0-255)

- Any TCP traffic (only) for a specific TCP port or range of ports, including optional control of connection traffic based on whether the initial request should be allowed
- Any UDP traffic (only) or UDP traffic for a specific UDP port
- Any ICMP traffic (only) or ICMP traffic of a specific type and code
- Any IGMP traffic (only) or IGMP traffic of a specific type
- Any of the above with specific precedence and/or ToS settings

Depending on the source and/or destination of a given IP traffic type, you must also determine the ACL application(s) (RACL, VACL, or static port ACL) needed to filter the traffic on the applicable switch interfaces. Answering the following questions can help you to design and properly position ACLs for optimum network usage.

- What are the logical points for minimizing unwanted IP traffic, and what ACL application(s) should be used? In many cases it makes sense to prevent unwanted IP traffic from reaching the core of your network by configuring ACLs to drop unwanted IP traffic at or close to the edge of the network. (The earlier in the network path you can block unwanted IP traffic, the greater the benefit for network performance.)
- From where is the traffic coming? The source and destination of IP traffic you want to filter determines the ACL application to use (RACL, VACL, static port ACL, and dynamic port ACL).
- What IP traffic should you explicitly block? Depending on your network size and the access requirements of individual hosts, this can involve creating a large number of ACEs in a given ACL (or a large number of ACLs), which increases the complexity of your solution.
- What IP traffic can you implicitly block by taking advantage of the implicit **deny IP any** to deny IP traffic that you have not explicitly permitted? This can reduce the number of entries needed in an ACL.
- What IP traffic should you permit? In some cases you will need to explicitly identify permitted IP traffic. In other cases, depending on your policies, you can insert an ACE with “permit any” forwarding at the end of an ACL. This means that all IP traffic not specifically matched by earlier entries in the list will be permitted.

Security

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

- blocking access from specific devices or interfaces (port or VLAN)
- blocking access to or from subnets in your network
- blocking access to or from the internet
- blocking access to sensitive data storage or restricted equipment
- preventing specific IP, TCP, UDP, IGMP, and ICMP traffic types, including unauthorized access using functions such as Telnet, SSH, and web browser

You can also enhance switch management security by using ACLs to block 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 in the switches covered by this guide do not filter non-IP traffic such as AppleTalk and IPX.

Guidelines for Planning the Structure of an ACL

After determining the filtering type (standard or extended) and ACL application (RACL, VACL, or static port ACL) to use at a particular point in your network, determine the order in which to apply individual ACEs to filter IP traffic (For information on ACL applications, refer to “ACL Applications” on page 10-15.) .

- The sequence of ACEs is significant. When the switch uses an ACL to determine whether to permit or deny a packet on a particular VLAN, it compares the packet to the criteria specified in the individual

Access Control Entries (ACEs) in the ACL, beginning with the first ACE in the list and proceeding sequentially until a match is found. When a match is found, the switch applies the indicated action (permit or deny) to the packet.

- The first match in an ACL dictates the action on a packet. Subsequent matches in the same ACL are ignored. However, if a packet is permitted by one ACL assigned to an interface, but denied by another ACL assigned to the same interface, the packet will be denied on the interface.
- On any ACL, the switch implicitly denies IP packets that are not explicitly permitted or denied by the ACEs configured in the ACL. If you want the switch to forward a packet for which there is not a match in an ACL, append an ACE that enables Permit Any forwarding as the last ACE in an ACL. This ensures that no packets reach the Implicit Deny case for that ACL.
- Generally, you should list ACEs from the most specific (individual hosts) to the most general (subnets or groups of subnets) unless doing so permits IP 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

- **RACLs and Routed IP Traffic:** Except for any IP traffic with a DA on the switch itself, RACLs filter only routed IP traffic that is entering or leaving the switch on a given VLAN. Thus, if routing is not enabled on the switch, there is no routed IP traffic for RACLs to filter. For more on routing, refer to the chapter titled “IP Routing Features” in the *Multicast and Routing Guide* for your switch.
- **VACLs and Switched or Routed IP Traffic:** A VACL filters any IP traffic entering the switch on the VLAN(s) to which it is assigned.
- **Static Port ACLs:** A static port ACL filters any IP traffic entering the switch on the port(s) or trunk(s) to which it is assigned.
- **Per Switch ACL Limits for All ACL Types.** At a minimum an ACL must have one, explicit “permit” or “deny” Access Control Entry. You can configure up to 2048 ACL assignments, as follows:
 - Named (Extended or Standard) ACLs: Up to 2048 (minus any numeric standard or extended ACL assignments)

- Numeric Standard ACLs: Up to 99; numeric range: 1 - 99
 - Numeric Extended ACLs: Up to 100; numeric range: 100 - 199
 - Total ACEs in all ACLs: Depends on the combined resource usage by ACL, QoS, IDM, Virus-Throttling, ICMP, and Management VLAN features (For more on this topic, refer to “Monitoring Shared Resources” on page 10-114.)
-
- **Implicit Deny:** 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 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 10-7 on page 10-29.
 - **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. Any ACEs listed after that point do not have any effect.
 - **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. Any ACEs listed after that point have no effect.
 - **Replacing One ACL with Another Using the Same Application:** For a specific interface, the most recent ACL assignment using a given application replaces any previous ACL assignment using the same application on the same interface. For example, if you configured an RACL named “100” to filter inbound routed IP traffic on VLAN 20, but later, you configured another RACL named 112 to filter inbound routed IP traffic on this same VLAN, RACL 112 replaces RACL 100 as the ACL to use.
 - **Static Port ACLs:** These are applied per-port, per port-list, or per static trunk. Adding a port to a trunk applies the trunk’s ACL configuration to the new member. If a port is configured with an ACL, the ACL must be removed before the port is added to the trunk. Also, removing a port from an ACL-configured trunk removes the ACL configuration from that port.

- **VACLs:** These filter any IP traffic entering the switch through any port belonging to the designated VLAN. VACLs do not filter IP traffic leaving the switch or being routed from another VLAN.
- **VACLs and RACLs Operate On Static VLANs:** You can assign an ACL to any VLAN that is statically configured on the switch. ACLs do not operate with dynamic VLANs.
- **A VACL or RACL Affects All Physical Ports in a Static VLAN:** A VACL or RACL assigned to a VLAN applies to all physical ports on the switch belonging to that VLAN, including ports that have dynamically joined the VLAN.
- **RACLs Screen Routed IP Traffic Entering or Leaving the Switch on a Given VLAN Interface:** This means that the following traffic is subject to ACL filtering:
 - IP traffic arriving on the switch through one VLAN and leaving the switch through another VLAN
 - IP traffic arriving on the switch through one subnet and leaving the switch through another subnet within the same, multinetted VLAN

Filtering the desired, routed IP traffic requires assigning an RACL to screen IP traffic inbound or outbound on the appropriate VLAN(s). In the case of a multinetted VLAN, it means that IP traffic inbound from different subnets in the same VLAN is screened by the same inbound RACL, and IP traffic outbound from different subnets is screened by the same outbound RACL. (Refer to figure 10-1 on page 10-17.)

- **RACLs Do Not Filter Switched IP Traffic Unless the Switch Itself is the SA or DA:** RACLs do *not* filter IP traffic moving between ports belonging to the same VLAN or subnet (in the case of a subnetted VLAN). (IP traffic moving between ports in different subnets of the same VLAN can be filtered.)

Note

RACLs do filter routed *or* switched IP traffic having an SA or DA on the switch itself.

How an ACE Uses a Mask To Screen Packets for Matches

When the switch applies an ACL to IP traffic, 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
10.38.252.195	255.255.255.0	first three octets	The fourth octet.
10.38.252.195	255.255.248.0	first two octets and the left-most five bits of the third octet	The right most three bits of the third octet and all bits in the fourth octet.

Thus, the bits set to 1 in a network mask define the part of an 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 criteria for determining whether to deny or permit a packet, or to pass it to the next ACE in the list. If there is a match, the configured deny or permit action occurs. If there is not a match, the packet is compared with the next ACE in the ACL. Thus, where a standard network mask defines how to identify the network and host numbers in an 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 10-39.

- 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 ACE.
- 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 10.28.100.15
```

produces this policy in an ACL listing:

IP Address	Mask
10.28.100.15	0.0.0.0

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

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

```
access-list 1 permit 10.28.32.1 0.0.0.31
```

IP Address	Mask
10.28.32.1	0.0.0.31

This policy states that:

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

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

Example of How the Mask Bit Settings Define a Match . Assume an ACE where the second octet of the mask for an SA is 7 (the rightmost three bits are “on”, or “1”) and the second octet of the corresponding SA in the ACE is 31 (the rightmost five bits). In this case, a match occurs when the second octet of the SA in a packet being filtered has a value in the range of 24 to 31. Refer to table 10-4, below.

Table 10-4. 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 10-8 to filter inbound packets on VLAN 20. Because the mask is all zeros, the ACE policy dictates that a match occurs only when the source IP address on such packets is identical to the IP address configured in the ACE.

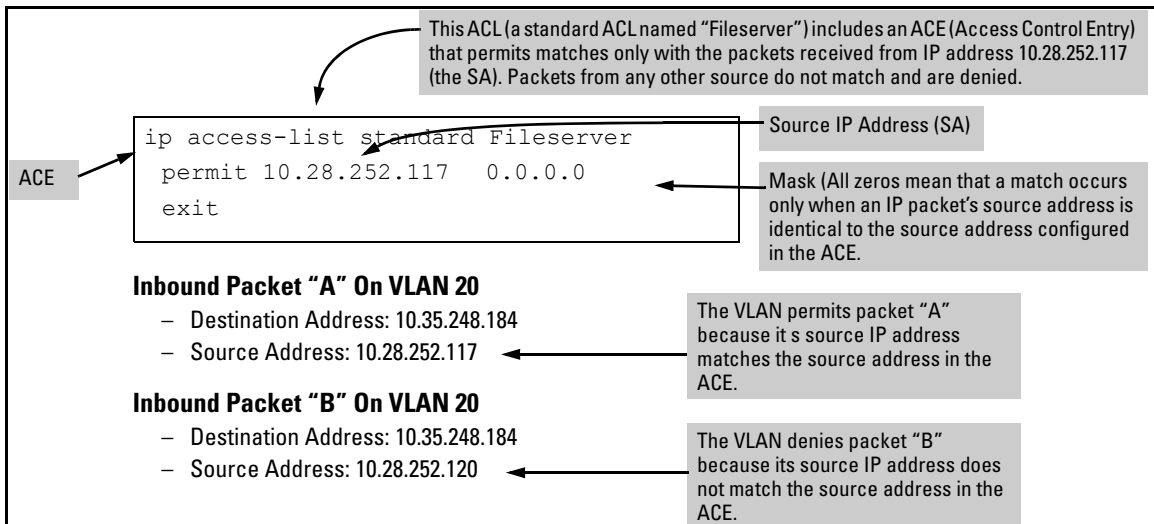


Figure 10-8. Example of an ACL with an Access Control Entry (ACE) that Allows Only One Source IP Address

Examples Allowing Multiple IP Addresses. Table 10-5 provides examples of how to apply masks to meet various filtering requirements.

Table 10-5. 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 10-6, 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 10-6, 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 10-6, 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 10-6, below.)

Table 10-6. Mask Effect on Selected Octets of the IP Addresses in Table 10-5

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 10-41.)

CIDR Notation. For information on using CIDR notation to specify ACL masks, refer to “Using CIDR Notation To Enter the ACL Mask” on page 10-50.

Configuring and Assigning an ACL

ACL Feature	Page
Configuring and Assigning a Standard ACL	10-51
Configuring and Assigning an Extended ACL	10-60
Enabling or Disabling ACL Filtering	10-81

Overview

General Steps for Implementing ACLs

1. Configure one or more ACLs. This creates and stores the ACL(s) in the switch configuration.
2. Assign an ACL. This step uses one of the following applications to assign the ACL to an interface:
 - RACL (routed IP traffic entering or leaving the switch on a given VLAN)
 - VACL (any IP traffic entering the switch on a given VLAN)
 - Static Port ACL (any IP traffic entering the switch on a given port, port list, or static trunk)
3. If the ACL is applied as an RACL, enable IP routing. Except for instances where the switch is the traffic source or destination, assigned RACLs filter IP traffic only when routing is enabled on the switch.

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**.

Options for Permit/Deny Policies

The permit or deny policy for IP traffic you want to filter can be based on source IP address alone, or on source IP address plus other IP factors.

- **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
 - IP protocol options:
 - Any IP traffic
 - Any IP traffic of a specific protocol type (0-255)
 - Any TCP traffic (only) for a specific TCP port or range of ports, including optional control of connection traffic based on whether the initial request should be allowed
 - Any UDP traffic (only) or UDP traffic for a specific UDP port
 - Any ICMP traffic (only) or ICMP traffic of a specific type and code
 - Any IGMP traffic (only) or IGMP traffic of a specific type
 - Any of the above with specific precedence and/or ToS settings

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

Carefully plan ACL applications before configuring specific ACLs. For more on this topic, refer to “Planning an ACL Application” on page 10-30.

ACL Configuration Structure

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

The basic ACL structure includes four elements:

1. ACL identity and type: This identifies the ACL as **standard** or **extended** and shows the ACL name or number.
2. Optional **remark** entries.

3. One or more deny/permit list entries (ACEs): One entry per line.

Element	Notes
Type	Standard or Extended
Identifier	<ul style="list-style-type: none"> • Alphanumeric; Up to 64 Characters, Including Spaces • Numeric: 1 - 99 (Standard) or 100 - 199 (Extended)
Remark	Allows up to 100 alphanumeric characters, including blank spaces. (If any spaces are used, the remark must be enclosed in a pair of single or double quotes.) A remark is associated with a particular ACE and will have the same sequence number as the ACE. (One remark is allowed per ACE.) Refer to "Attaching a Remark to an ACE" on page 10-92.
Maximum ACEs Per per Switch	The upper limit on ACEs supported by the switch depends on the concurrent resource usage by configured QoS, ICMP rate-limiting, management VLAN, and virus-throttling features. Refer to "Monitoring Shared Resources" on page 10-114.

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

Standard ACL Structure

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

```
ip access-list standard < identifier >
  [ [ seq-# ] remark < remark-str > ]
  < permit | deny > < SA > [ log ]
  .
  .
  .
  < Implicit Deny >
  exit
```

Note: The optional **log** function is available only for explicit "deny" ACEs.

Figure 10-9. Example of the General Structure for a Standard ACL

Access Control Lists (ACLs)

Configuring and Assigning an ACL

For example, figure 10-10 shows how to interpret the entries in a standard ACL.

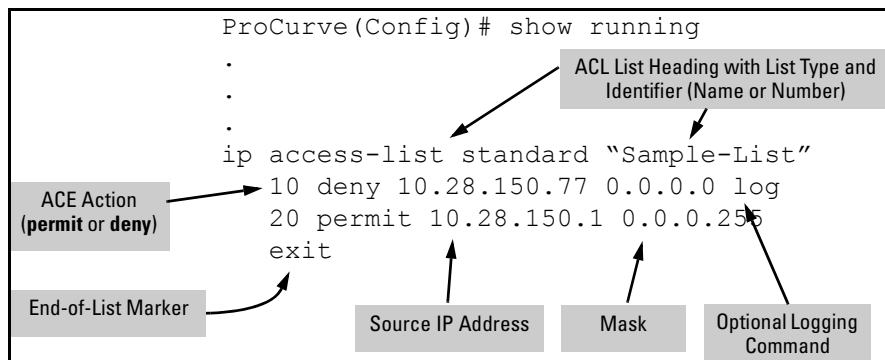


Figure 10-10. Example of a Displayed Standard ACL Configuration with Two ACEs

Extended ACL Configuration Structure

Individual ACEs in an extended ACL include:

- A permit/deny statement
- Source and destination IP addressing
- Choice of IP criteria, including optional precedence and ToS
- Optional ACL **log** command (for **deny** entries)
- Optional remark statements

```
ip access-list extended < identifier >
[[ seq-#] remark < remark-str >]
< permit | deny > < ip-type > < SA > < src-acl-mask > < DA > < dest-acl-mask > [log]

    < permit | deny > tcp
        < SA > < src-acl-mask > [< operator > < port-id >]
        < DA > < desti-acl-mask > [< operator > < port-id >] [log]
        [ established ]

    < permit | deny > udp
        < SA > < src-acl-mask > [< operator > < port-id >]
        < DA > < dest-acl-mask > [< operator > < port-id >] [log]

    < permit | deny > icmp
        < SA > < src-acl-mask > < DA > < dest-acl-mask > [ icmp-type ] [log]

    < permit | deny > igmp
        < SA > < SA-mask > < DA > < dest-acl-mask > [ igmp-type ] [log]

    [ precedence < priority >]
    [ tos < tos-setting >]
    . . .
    < Implicit Deny >
    exit
```

Note: The optional **log** function appears only with “deny” ACEs.

Figure 10-11. Example of General Structure Options for an Extended ACL

Access Control Lists (ACLs) Configuring and Assigning an ACL

For example, figure 10-12 shows how to interpret the entries in an extended ACL.

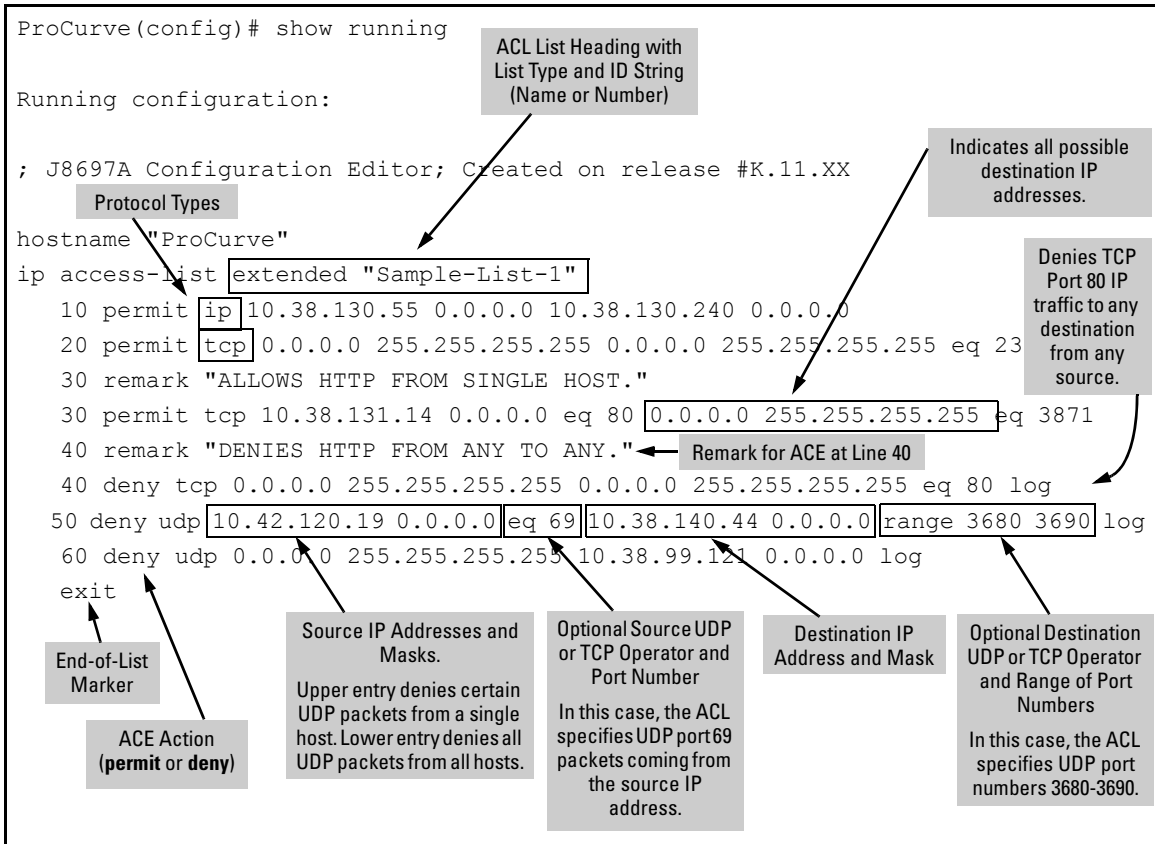


Figure 10-12. Example of a Displayed Extended ACL Configuration

ACL Configuration Factors

The Sequence of Entries in an ACL Is Significant

When the switch uses an ACL to determine whether to permit or deny a packet, it compares the packet to the criteria specified in the individual Access Control Entries (ACEs) in the ACL, beginning with the first ACE in the list and proceeding sequentially until a match is found. When a match is found, the switch applies the indicated action (permit or deny) to the packet. This is

significant because, once a match is found for a packet, subsequent ACEs in the same ACL will not be applied to that packet, regardless of whether they match the packet.

For example, suppose that you have applied the ACL shown in figure 10-13 to inbound IP traffic on VLAN 1 (the default VLAN):

```

ip access-list extended "Sample-List-2"
 10 deny ip 10.28.235.10 0.0.0.0 0.0.0.0 255.255.255.255
 20 deny ip 10.28.245.89 0.0.0.0 0.0.0.0 255.255.255.255
 30 permit tcp 10.28.18.100 0.0.0.0 10.28.237.1 0.0.0.0
 40 deny tcp 10.28.18.100 0.0.0.0 0.0.0.0 255.255.255.255
 50 permit ip 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255

(implicit deny) ←
exit

```

Figure 10-13. Example of a Standard ACL that Permits All IP Traffic Not Implicitly Denied

Table 10-7. Effect of the Above ACL on Inbound IP Traffic in the Assigned VLAN

Line #	Action
n/a	Shows type (extended) and ID (Sample-List-2).
10	A packet from IP source address 10.28.235.10 will be denied (dropped). This ACE filters out all packets received from 10.28.235.10. As a result, IP traffic from that device will not be allowed and packets from that device will not be compared against any later entries in the list.
20	A packet from IP source 10.28.245.89 will be denied (dropped). This ACE filters out all packets received from 10.28.245.89. As the result, IP traffic from that device will not be allowed and packets from that device will not be compared against any later entries in the list.
30	A TCP packet from SA 10.28.18.100 with a DA of 10.28.237.1 will be permitted (forwarded). Since no earlier ACEs in the list have filtered TCP packets from 10.28.18.100 and destined for 10.28.237.1, the switch will use this ACE to evaluate such packets. Any packets that meet this criteria will be forwarded. (Any packets that do not meet this TCP source-destination criteria are not affected by this ACE.)
40	A TCP packet from source address 10.28.18.100 to any destination address will be denied (dropped). Since, in this example, the intent is to block TCP traffic from 10.28.18.100 to any destination except the destination stated in the ACE at line 30, this ACE must follow the ACE at line 30. (If their relative positions were exchanged, all TCP traffic from 10.28.18.100 would be dropped, including the traffic for the 10.28.18.1 destination.)
50	Any packet from any IP source address to any destination address will be permitted (forwarded). The only traffic to reach this ACE will be IP packets not specifically permitted or denied by the earlier ACEs.

Line #	Action
<i>n/a</i>	The <i>Implicit Deny</i> is a function the switch automatically adds 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 ACL. In this example, the ACE at line 50 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 function.
exit	Marks the end of the ACL.

Allowing for the Implied Deny Function

In any ACL having one or more ACEs there will always be a packet match. This is because the switch automatically applies an Implicit Deny as the last ACE in any ACL. This function is not visible in ACL listings, but is always present. (Refer to figure 10-13.) This means that if you configure the switch to use an ACL for filtering either inbound or outbound IP traffic on a VLAN, any packets not specifically permitted or denied by the explicit entries you create will be denied by the Implicit Deny action. If you want to preempt the Implicit Deny (so that IP traffic not specifically addressed by earlier ACEs in a given ACL will be permitted), insert an explicit **permit any** (for standard ACLs) or **permit ip any any** (for extended ACLs) as the last explicit ACE in the ACL.

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

The switch stores ACLs in the configuration file. Thus, until you actually assign an ACL to an interface, it is present in the configuration, but not used (and does not use any of the monitored resources described in the appendix titled “Monitored Resources” in the *Management and Configuration Guide* for your switch.)

You Can Assign an ACL Name or Number to an Interface Even if the ACL Does Not Exist in the Switch’s Configuration

In this case, if you subsequently create an ACL with that name or number, the switch automatically applies each ACE as soon as you enter it in the running-config file. Similarly, if you modify an existing ACE in an ACL you already applied to an interface, the switch automatically implements the new ACE as soon as you enter it. (See “General ACL Operating Notes” on page 10-113.) The switch allows a maximum of 2048 ACLs in any combination of numeric and alphanumeric names, and determines the total from the number of unique ACL names in the configuration. For example, if you configure two ACLs, but assign only one of them to a VLAN, the ACL total is two, for the two unique ACL names. If you then assign the name of a nonexistent ACL to a VLAN, the new ACL total is three, because the switch now has three unique ACL names in its configuration.

Using the CLI To Create an ACL

Command	Page
access-list (standard ACLs)	10-51
access-list (extended ACLs)	10-60

You can use either the switch CLI or an offline text editor to create an ACL. This section describes the CLI method, which is recommended for creating short ACLs. (To use the offline method, refer to “Creating or Editing ACLs Offline” on page 10-104.)

General ACE Rules

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

- Inserting or adding an ACE to an ACL:
 - **Named ACLs:** Add an ACE to the end of a named ACE by using the **ip access-list** command to enter the Named ACL (**nacl**) context and entering the ACE without the sequence number. For example, if you wanted to add a “permit” ACL at the end of a list named “List-1” to allow IP traffic from the device at 10.10.10.100:

```
ProCurve(config)# ip access-list standard List-1
ProCurve(config-std-nacl)# permit host
10.10.10.100
```

Insert an ACE anywhere in a named ACL by specifying a sequence number. For example, if you wanted to insert a new ACE as line 15 between lines 10 and 20 in an existing ACL named “List-2” to deny IP traffic from the device at 10.10.10.77:

```
ProCurve(config)# ip access-list standard List-2
ProCurve(config-std-nacl)# 15 deny host 10.10.10.77
```

- **Numbered ACLs:** Add an ACE to the end of a numbered ACL by using the **access-list <1-99|100-199>** command. For example, if you wanted to add a “permit” ACE at the end of a list identified with the number “11” to allow IP traffic from the device at 10.10.10.100:

```
ProCurve(config)# access-list 11 permit host
10.10.10.100
```

To insert an ACE *anywhere* in a numbered ACL, use the same process as described above for inserting an ACE anywhere in a *named* ACL. For example, to insert an ACE denying IP traffic from the host at 10.10.10.77 as line 52 in an existing ACL identified (named) with the number 11:

```
ProCurve(config)# ip access-list standard 99
ProCurve(config-std-nacl)# 52 deny host 10.10.10.77
```

Note

After a numbered ACL has been created (using **access-list < 1 - 99 | 100 - 199 >**), it can be managed as either a named or numbered ACL, as shown above.

- Deleting an ACE: Enter the ACL context and delete the sequence number for the unwanted ACE. (To view the sequence numbers of the ACEs in a list, use **show access-list < acl-name-str >**.)
- Duplicate ACEs are not allowed in the same ACL. Attempting to enter a duplicate ACE displays the **Duplicate access control entry** message.

Using CIDR Notation To Enter the 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 that must match. The switch then converts the mask to inverse notation for ACL use.

Table 10-8. Examples of CIDR Notation for Masks

IP Address Used In an ACL with CIDR Notation	Resulting ACL Mask	Meaning
10.38.240.125/15	0.1.255.255	The leftmost 15 bits must match; the remaining bits are wildcards.
10.38.240.125/20	0.0.15.255	The leftmost 20 bits must match; the remaining bits are wildcards.
10.38.240.125/21	0.0.7.255	The leftmost 21 bits must match; the remaining bits are wildcards.
10.38.240.125/24	0.0.0.255	The leftmost 24 bits must match; the remaining bits are wildcards.
18.38.240.125/32	0.0.0.0	All bits must match.

Configuring Standard ACLs

Table 10-9. Command Summary for Standard ACLs

Action	Command(s)	Page
Create a Standard, Named ACL	ProCurve(config)# ip access-list standard < name-str > ProCurve(config-std-nacl)# < deny permit >	10-53
or Add an ACE to the End of an Existing Standard, Named ACL	< any host <SA > SA/< mask-length > SA < mask >> ¹ [log] ²	
Create a Standard, Numbered ACL	ProCurve(config)# access-list < 1-99 > < deny permit > < any host <SA > SA/< mask-length > SA < mask >> [log] ²	10-56
or Add an ACE to the End of an Existing Standard, Numbered ACL		
Use a Sequence Number To Insert an ACE in an ACL	ProCurve(config)# ip access-list standard < name-str 1-99 > ProCurve(config-std-nacl)# 1-2147483647 < deny permit > < any host <SA > SA/< mask-length > SA < mask >> ¹ [log] ²	10-87
Use an ACE's Sequence Number To Delete the ACE from an ACL	ProCurve(config)# ip access-list standard < name-str 1-99 > ProCurve(config-std-nacl)# no < 1-2147483647 >	10-90
Resequence the ACEs in an ACL	ProCurve(config)# ip access-list resequence < name-str 1-99 > < 1-2147483646 >	10-91
Enter or Remove a Remark from an ACL	ProCurve(config)# ip access-list standard < name-str 1-99 > ProCurve(config-ext-nacl)# [remark < remark-str > no < 1-2147483647 > remark]	10-92 10-94
<i>For numbered, standard ACLs only, the following remark commands can be substituted for the above:</i>		
	ProCurve(config)# access-list < 1 - 99 > remark < remark-str > ProCurve(config)# [no] access-list < 1 - 99 > remark	
Delete an ACL	ProCurve(config)# no ip access-list standard < name-str 1-99 >	10-85
<i>For numbered, standard ACLs, the following command can be substituted for the above:</i>		
	ProCurve(config)# access-list < 1 - 99 > remark < remark-str >	
<p>¹The mask can be in either dotted-decimal notation (such as 0.0.15.255) or CIDR notation (such as /20). ²The [log] function applies only to “deny” ACLs, and generates a message only when there is a “deny” match.</p>		

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

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

A *named*, standard ACL is identified by an alphanumeric string of up to 64 characters and is created by entering the Named ACL (**nacl**) context. A *numbered*, standard ACL is identified by a number in the range of 1 - 99 and is created without having to leave the global config context. Note that the CLI command syntax for creating a named ACL differs from the command syntax for creating a numbered ACL. For example, the first pair of entries below illustrate how to create (or enter) a named, standard ACL and enter an ACE. The next entry illustrates creating a numbered, standard ACL with the same ACE.

```
ProCurve(config)# ip access-list standard Test-List
ProCurve(config-std-nacl)# permit host 10.10.10.147
```

```
ProCurve(config)# access-list 1 permit host 10.10.10.147
```

Note that once a numbered ACL has been created, it can be accessed using the named ACL method. This is useful if it becomes necessary to edit a numbered ACL by inserting or removing individual ACEs. (Inserting or deleting an ACE is done by sequence number, and requires the Named ACL (**nacl**) context.) The switch allows a maximum of 2048 unique ACL identities; standard and extended combined.

Note

For a summary of standard ACL commands, refer to table 10-9 on page 10-51. For a summary of all ACL commands, refer to tables 10-1 and 10-2 on pages 10-6 and 10-8.

Configuring Named, Standard ACLs

This section describes the commands for performing the following:

- creating and/or entering the context of a named, standard ACL
- appending an ACE to the end of an existing list or entering the first ACE in a new list

For other ACL topics, refer to the following:

Topic	Page
configuring numbered, standard ACLs	10-56
configuring named, extended ACLs	10-62
configuring numbered, extended ACLs	10-74
applying or removing an ACL on an interface	10-81
deleting an ACL	10-85
editing an ACL	10-86
sequence numbering in ACLs	10-87
including remarks in an ACL	10-92
displaying ACL configuration data	10-96
creating or editing ACLs offline	10-104
enabling ACL "Deny" logging	10-109

Entering the "Named ACL" (nacl) Context. This command is a prerequisite to entering or editing ACEs in a named ACL.

Syntax: ip access-list standard < name-str >

Places the CLI in the "Named ACL" (nacl) context specified by the < name-str > alphanumeric identifier. This enables entry of individual ACEs in the specified ACL. If the ACL does not already exist, this command creates it.

< name-str >: *Specifies an identifier for the ACL. Consists of an alphanumeric string of up to 64 case-sensitive characters. Including spaces in the string requires that you enclose the string in single or double quotes. For example: "Accounting ACL".*

Refer also to table 10-9 on page 10-51.

Configuring ACEs in a Named, Standard ACL. Configuring ACEs is done after using the **ip access-list standard < name-str >** command described above to enter the “Named ACL” (**nacl**) context of an access list. *For a standard ACL syntax summary, refer to table 10-9 on page 10-51.*

Syntax: < deny | permit >
< any | host < SA > | SA < mask | SA / mask-length >> [log]

*Executing this command appends the ACE to the end of the list of ACEs in the current ACL. In the default ACL configuration, ACEs are automatically assigned consecutive sequence numbers in increments of 10 and can be renumbered using **resequence** (page 10-91).*

Note: *To insert a new ACE between two existing ACEs, precede **deny** or **permit** with an appropriate sequence number. (Refer to “Inserting an ACE in an Existing ACL” on page 10-88.)*

< deny | permit >

*For named ACLs, used in the “Named ACL” (**nacl**) context to configure an ACE. Specifies whether the ACE denies or permits a packet matching the criteria in the ACE, as described below.*

< any | host < SA > | SA < mask > | SA / mask-length >

Defines the source IP address (SA) a packet must carry for a match with the ACE.

- **any** — Allows IP packets from any SA.
- **host < SA >** — Specifies only packets having < SA > as the source. Use this criterion when you want to match the IP packets from a single source IP address.
- **SA < mask >** or **SA / mask-length** — Specifies packets received from either a subnet or a group of IP addresses. The mask format can be in either dotted-decimal format or CIDR format (number of significant bits). (Refer to “Using CIDR Notation To Enter the ACL Mask” on page 10-50).

Mask Application: *The mask is applied to the IP address in the ACE to define which bits in a packet’s source IP address must exactly match the IP address configured in the ACE and which bits need not match. For example: **10.10.10.1/24** and **10.10.10.1 0.0.0.255** both define any IP address in the range of 10.10.10.(1 - 255).*

Note: *Specifying a group of contiguous IP addresses may require more than one ACE. For more on how masks operate, refer to “How an ACE Uses a Mask To Screen Packets for Matches” on page 10-36.*

[log]

This option 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 10-109.)*

(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 10-109.)

Example of Creating and Listing a Standard, Named ACL. This example illustrates how to create a standard, named ACL with several ACEs. This example creates an ACL that:

1. permits IP traffic from a host with the IP address of 10.10.10.104
2. creates another ACE that blocks all other IP traffic from the same subnet
3. allows all other IP traffic

<pre>ProCurve(config)# ip access-list standard Sample-List ProCurve(config-std-nacl)# permit host 10.10.10.104 ProCurve(config-std-nacl)# deny 10.10.10.1/24 log ProCurve(config-std-nacl)# permit any ProCurve(config-std-nacl)# exit ProCurve(config)# _</pre>	<p>Creates the “Sample-List” ACL and enters the “Named ACL” context for this list.</p> <p>Appends three ACEs to the list in the order shown.</p> <p>Exits from the nacl context.</p>
--	--

Figure 10-14. Example of Commands Used To Create a Standard, Named ACL

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

Access Control Lists

Name: Sample-List
Type: Standard
Applied: No

SEQ  Entry
-----
10   Action: permit
     IP      : 10.10.10.104      Mask: 0.0.0.0
20   Action: deny (log)
     IP      : 10.10.10.1       Mask: 0.0.0.255
30   Action: permit
     IP      : 0.0.0.0          Mask: 255.255.255.255
```

Note that each ACE is automatically assigned a sequence number.

Figure 10-15. Screen Output Listing the “Sample-List” ACL Content

Creating Numbered, Standard ACLs

Use the following general steps to create or add to a numbered, standard ACL:

1. Create a numbered, standard ACL by entering the first ACE in the list
2. Append a new ACE to the end of an existing, standard ACL

This section describes the commands for performing these steps. For other ACL topics, refer to the following:

Topic	Page
configuring named, standard ACLs	10-53
configuring named, extended ACLs	10-62
configuring numbered, extended ACLs	10-74
applying or removing an ACL on an interface	10-81
deleting an ACL	10-85
editing an ACL	10-86
sequence numbering in ACLs	10-87
including remarks in an ACL	10-92
displaying ACL configuration data	10-96
creating or editing ACLs offline	10-104
enabling ACL “Deny” logging	10-109

Creating or Adding to a Standard, Numbered ACL. *This command is an alternative to using **ip access-list standard < name-str >** and does not use the “Named ACL” (**nacl**) context. For a standard ACL syntax summary, refer to table 10-9 on page 10-51.*

Syntax: access-list < 1-99 > < deny | permit >
< any | host < SA > | SA < mask | SA/mask-length >> [log]

*Appends an ACE to the end of the list of ACEs in the current standard, numbered ACL. If the ACL does not already exist, creates both the ACL and its first ACE. In the default configuration, ACEs are automatically assigned consecutive sequence numbers in increments of 10 and can be renumbered using **resequence** (page 10-91).*

Note: *To insert a new ACE between two existing ACEs in a standard, numbered ACL:*

- a. *Use **ip access list extended < 1 - 99 >** to open the ACL as a named ACL.*
- b. *Enter the desired sequence number along with the ACE keywords and variables you want.*

(After a numbered ACL has been created, it can be managed as either a named or numbered ACL. Refer to the “Numbered ACLs” list item on page 10-49.)

< 1-99 >

*Specifies the ACL identifier as a number. The switch interprets an ACL with a value in this range as a standard ACL (which filters all IP traffic on the basis of SA). (To create a standard access list with an alphanumeric name (**name-str**) instead of a number, refer to “Configuring Named, Standard ACLs” on page 10-53.)*

< deny | permit >

Specifies whether the ACE denies or permits a packet matching the criteria in the ACE, as described next.

< any | host < SA > | SA < mask | SA/mask-length >>

Defines the source IP address (SA) a packet must carry for a match with the ACE.

- **any** — Allows IP packets from any SA.
- **host < SA >** — Specifies only packets having < SA > as the source. Use this criterion when you want to match only the IP packets from a single SA.
- **SA < mask > or SA /mask-length** — Specifies packets received from an SA, where the SA is either a subnet or a group of IP addresses. The mask format can be in either dotted-decimal format or CIDR format (number of significant bits). (Refer to “Using CIDR Notation To Enter the ACL Mask” on page 10-50).

SA Mask Application: *The mask is applied to the SA in the ACE to define which bits in a packet’s SA must exactly match the SA configured in the ACL and which bits need not match.*

Example: **10.10.10.1/24** and **10.10.10.1 0.0.0.255** both define any IP address in the range of 10.10.10.(1 - 255).

Note: *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 10-36.*

Example of Creating and Viewing a Standard ACL. This example creates a standard, numbered ACL with the same ACE content as show in figure 10-14 on page 10-55.

```
ProCurve(config)# access-list 17 permit host 10.10.10.104
ProCurve(config)# access-list 17 deny 10.10.10.1/24 log
ProCurve(config)# access-list 17 permit any
ProCurve(config)# show access-list 17
```

Access Control Lists

Name: 17
Type: Standard
Applied: No

SEQ Entry

```
-----
10  Action: permit
    IP      : 10.10.10.104      Mask: 0.0.0.0

20  Action: deny (log)
    IP      : 10.10.10.1       Mask: 0.0.0.255

30  Action: permit
    IP      : 0.0.0.0          Mask: 255.255.255.255
```

Note that each ACE is automatically assigned a sequence number.

Figure 10-16. Standard, Numbered ACL with the Same ACEs as the Standard, Named ACL in Figure 10-14

Configuring Extended ACLs

Table 10-10. Command Summary for Extended ACLs

Action	Command(s)	Page
Create an Extended, Named ACL <i>or</i> Add an ACE to the End of an Existing, Extended ACL	<pre> ProCurve(config)# ip access-list extended < name-str 100-199 > ProCurve(config-std-nacl)# < deny permit > < ip ip-protocol ip-protocol-nbr > < any host <SA> SA/< mask-length > SA < mask >>¹ < any host <DA> DA/< mask-length > DA < mask >>¹ [tcp udp] < any host <SA> SA/< mask-length > SA < mask >>¹ [comparison-operator < value >]] < any host <DA> DA/< mask-length > DA < mask >>¹ [comparison-operator < value >] [established] < igmp > < any host <SA> SA/< mask-length > SA < mask >>¹ < any host <DA> DA/< mask-length > DA < mask >>¹ [igmp-packet-type] < icmp > < any host <SA> SA/< mask-length > SA < mask >>¹ < any host <DA> DA/< mask-length > DA < mask >>¹ [[< 0 - 255 > [0 - 255]] icmp-message] [precedence < priority >] [tos < tos- setting >] [log]² </pre>	10-62
Create an Extended, Numbered ACL <i>or</i> Add an ACE to the End of an Existing, Numbered ACL	<pre> ProCurve(config)# access-list < 100-199 > < deny permit > < ip-options tcp/udp-options igmp-options icmp-options > [log]² [precedence < priority >] [tos < tos- setting >] </pre> <p>Note: Uses the same IP, TCP/UDP, IGMP, and ICMP options as shown above for "Create an Extended, Named ACL".</p>	10-74
Insert an ACE by Assigning a Sequence Number	<pre> ProCurve(config)# ip access-list extended < name-str 100-199 > ProCurve(config-ext-nacl)# 1-2147483647 < deny permit > </pre> <p><i>Uses the options shown above for "Create an Extended, Named ACL".</i></p>	10-88
Use Sequence Number To Delete an ACE	<pre> ProCurve(config)# ip access-list extended < name-str 100-199 > ProCurve(config-std-nacl)# no < 1-2147483647 > </pre>	10-90
Resequence the ACEs in an ACL	<pre> ProCurve(config)# ip access-list resequence < name-str 100-199 > < 1-2147483647 > < 1-2147483646 > </pre>	10-91

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

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

Table continues on the next page.

Action	Command(s)	Page
Enter or Remove a Remark	ProCurve(config)# ip access-list extended < name-str 100-199 >	10-92
	ProCurve(config-ext-nacl)# [remark < remark-str > no < 1 - 2147483647 > remark]	10-94
	<i>For numbered, extended ACLs only, the following remark commands can be substituted for the above:</i>	
	ProCurve(config)# access-list < 100 - 199 > remark < remark-str >	
	ProCurve(config)# [no] access-list < 100 - 199 > remark	
Delete an Extended ACL	ProCurve(config)# no ip access-list extended < name-str 100-199 >	10-85
	<i>For numbered, extended ACLs only, the following command can also be used:</i>	
	ProCurve(config)# no access-list < 100 - 199 >	

Standard ACLs use only source IP addresses for filtering criteria, extended ACLs use multiple filtering criteria. This enables you to more closely define your IP packet-filtering. Extended ACLs enable filtering on the following:

- Source and destination IP addresses (required), in one of the following options:
 - specific host IP
 - subnet or group of IP addresses
 - any IP address
- choice of any IP protocol
- optional packet-type criteria for IGMP, and ICMP traffic
- optional source and/or destination TCP or UDP port, with a further option for comparison operators and (for TCP) an option for establishing connections
- filtering for TCP traffic based on whether the subject traffic is initiating a connection (“established” option)
- optional IP precedence and ToS criteria

The switch allows up to 2048 ACLs in any combination of numeric and alphanumeric identifiers, and determines the total from the number of unique identifiers in the configuration. For example, configuring two ACLs results in an ACL total of two, even if neither is assigned to an interface. If you then assign a nonexistent ACL to an interface, the new ACL total is three, because the switch now has three unique ACL names in its configuration. (For more on ACL limits, refer to “Monitoring Shared Resources” on page 10-114.)

Configuring Named, Extended ACLs

For a match to occur with an ACE in an extended ACL, a packet must have the source and destination IP address criteria specified by the ACE, as well as any IP protocol-specific criteria included in the command.

Use the following general steps to create or add to a named, extended ACL:

1. Create and/or enter the context of a named, extended ACL.
2. Enter the first ACE in a new, extended ACL or append an ACE to the end of an existing, extended ACL.

This section describes the commands for performing these steps. For other ACL topics, refer to the following:

Topic	Page
configuring named, standard ACLs	10-53
configuring numbered, standard ACLs	10-56
configuring numbered, extended ACLs	10-74
applying or removing an ACL on an interface	10-81
deleting an ACL	10-85
editing an ACL	10-86
sequence numbering in ACLs	10-87
including remarks in an ACL	10-92
displaying ACL configuration data	10-96
creating or editing ACLs offline	10-104
enabling ACL "Deny" logging	10-109

Creating a Named, Extended ACL and/or Entering the “Named ACL” (nacl) Context. This command is a prerequisite to entering or editing ACEs in a named, extended ACL. (For a summary of the extended ACL syntax options, refer to table 10-10 on page 10-60.)

Syntax: ip access-list extended < name-str >

Places the CLI in the “Named ACL” (nacl) context specified by the < name-str > alphanumeric identifier. This enables entry of individual ACEs in the specified ACL. If the ACL does not already exist, this command creates it.

< name-str >: *Specifies an alphanumeric identifier for the ACL. Consists of an alphanumeric string of up to 64 case-sensitive characters. Including spaces in the string requires that you enclose the string in single or double quotes. For example: “Accounting ACL”. You can also use this command to access an existing, numbered ACL. Refer to “Using the CLI To Edit ACLs” on page 10-86*

```
ProCurve(config)# ip access-list extended Sample-List
ProCurve(config-ext-nacl)#
```

Figure 10-17. Example of Entering the Named ACL Context

Configure ACEs in a Named, Extended ACL and/or Enter the “Named ACL” (nacl) Context. Configuring ACEs is done after using the **ip access-list standard < name-str >** command described on page 10-63 to enter the “Named ACL” (**nacl**) context of an ACL. For an extended ACL syntax summary, refer to table 10-10 on page 10-60.

Syntax: < deny | permit > < ip | ip-protocol | ip-protocol-nbr >
(nacl
context) < any | host < SA > | SA / mask-length | SA < mask > >
< any | host < DA > | DA / mask-length | DA < mask > >
[precedence] [tos] [log]

*Appends an ACE to the end of the list of ACEs in the current ACL. In the default configuration, ACEs are automatically assigned consecutive sequence numbers in increments of 10 and can be renumbered using **resequence** (page 10-91).*

Note: *To insert a new ACE between two existing ACEs in an extended, named ACL, precede **deny** or **permit** with an appropriate sequence number along with the ACE keywords and variables you want. (Refer to “Inserting an ACE in an Existing ACL” on page 10-88.)*

For a match to occur, a packet must have the source and destination IP addressing criteria specified in the ACE, as well as:

- *the protocol-specific criteria configured in the ACE, including any included, optional elements (described later in this section)*
- *any (optional) precedence and/or ToS settings configured in the ACE*

< deny | permit >

*For named ACLs, these keywords are used in the “Named ACL” (**nacl**) context to specify whether the ACE denies or permits a packet matching the criteria in the ACE, as described below.*

< ip | ip-protocol | ip-protocol-nbr >

Used after **deny** or **permit** to specify the packet protocol type required for a match. An extended ACL must include one of the following:

- **ip** — any IP packet.
- **ip-protocol** — any one of the following IP protocol names:

ip-in-ip	ipv6-in-ip	gre	esp	ah
ospf	pim	vrrp	sctp	tcp*
udp*	icmp*	igmp*		
- **ip-protocol-nbr** — the IPv4 IP protocol number of an IP packet type, such as “8” for Exterior Gateway Protocol or 121 for Simple Message Protocol. (For a listing of IP protocol numbers and their corresponding protocol names, refer to the IANA “Protocol Number Assignment Services” at www.iana.com.) (Range: 0 - 255)

* For TCP, UDP, ICMP, and IGMP, additional criteria can be specified, as described on pages 10-68 through 10-72.

< any | host < SA > | SA < mask > | SA / mask-length

This is the first instance of IP addressing in an extended ACE. It follows the protocol specifier and defines the source IP address (SA) a packet must carry for a match with the ACE.

- **any** — Allows IP packets from any SA.
- **host < SA >** — Specifies only packets having a single address as the SA. Use this criterion when you want to match only the IP packets from a single SA.
- **SA < mask >** or **SA / mask-length** — Specifies packets received from an SA, where the SA is either a subnet or a group of IP addresses. The mask can be in either dotted-decimal format or CIDR format (number of significant bits). Refer to “Using CIDR Notation To Enter the ACL Mask” on page 10-50.

SA Mask Application: The mask is applied to the SA in the ACL to define which bits in a packet’s SA must exactly match the SA configured in the ACL and which bits need not match.

Example: 10.10.10.1/24 and 10.10.10.1 0.0.0.255 both define any IP address in the range of 10.10.10.(1 - 255).

Note: 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 10-36.

< any | host < DA > | DA/mask-length | DA/ < mask >>

This is the second instance of IP addressing in an extended ACE. It follows the first (SA) instance, described earlier, and defines the destination IP address (DA) that a packet must carry in order to have a match with the ACE.

- **any** — Allows routed IP packets to any DA.
- **host < DA >** — Specifies only packets having **DA** as the destination address. Use this criterion when you want to match only the IP packets for a single DA.
- **DA/mask-length** or **DA < mask >** — Specifies packets intended for a destination address, where the address is either a subnet or a group of IP addresses. The mask format can be in either dotted-decimal format or CIDR format (number of significant bits). Refer to “Using CIDR Notation To Enter the ACL Mask” on page 10-50.

DA Mask Application: *The mask is applied to the DA in the ACL to define which bits in a packet’s DA must exactly match the DA configured in the ACL and which bits need not match. See also the above example and note.*

[precedence < 0 - 7 | precedence-name >]

This option can be used after the DA to cause the ACE to match packets with the specified IP precedence value. Values can be entered as the following IP precedence numbers or alphanumeric names:

0	or	routine
1	“	priority
2	“	immediate
3	“	flash
4	“	flash-override
5	“	critical
6	“	internet (for internetwork control)
7	“	network (for network control)

Note: *The precedence criteria described in this section are applied in addition to any other selection criteria configured in the same ACE.*

[tos < tos-setting >]

This option can be used after the DA to cause the ACE to match packets with the specified IP Type-of-Service (ToS) setting. ToS values can be entered as the following numeric settings or, in the case of 0, 2, 4, and 8, as alphanumeric names:

0	or	normal
2	“	max-reliability
4	“	max-throughput
6		
8	“	minimize-delay
10		
12		
14		

Note: *The ToS criteria in this section are applied in addition to any other criteria configured in the same ACE.*

[log]

This option can be used after the DA to generate an Event Log message if:

- *The action is **deny**. (Not applicable to **permit**.)*
- *There is a match.*
- *ACL logging is enabled. (Refer to “Enabling ACL Logging on the Switch” on page 10-111.)*

Options for TCP and UDP Traffic in Extended ACLs. An ACE designed to permit or deny TCP or UDP traffic can optionally include port number criteria for either the source or destination, or both. Use of TCP criteria also allows the **established** option for controlling TCP connection traffic. (For a summary of the extended ACL syntax options, refer to table 10-10 on page 10-60.)

Syntax: < deny | permit > < tcp | udp >
< SA > [comparison-operator < tcp/udp-src-port >]
< DA >
[comparison-operator < tcp-dest-port >] [established]
[comparison-operator < udp-dest-port >]

*In an extended ACL using either **tcp** or **udp** as the IP packet protocol type, you can optionally use TCP or UDP source and/or destination port numbers or ranges of numbers to further define the criteria for a match. For example:*

```
#deny tcp host 10.20.10.17 eq 23 host 10.20.10.155
  established
#permit tcp host 10.10.10.100 host 10.20.10.17
  eq telnet
#deny udp 10.30.10.1/24 host 10.20.10.17 range
  161 162
```

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

To specify a TCP or UDP source port number in an ACE, (1) select a comparison operator from the following list and (2) enter the port number or a well-known port name.

Comparison Operators:

- **eq** < **tcp/udp-port-nbr** > — “Equal To”; to have a match with the ACE entry, the TCP or UDP source port number in a packet must be equal to < **tcp/udp-port-nbr** >.
- **gt** < **tcp/udp-port-nbr** > — “Greater Than”; to have a match with the ACE entry, the TCP or UDP source port number in a packet must be greater than < **tcp/udp-port-nbr** >.
- **lt** < **tcp/udp-port-nbr** > — “Less Than”; to have a match with the ACE entry, the TCP or UDP source port number in a packet must be less than < **tcp/udp-port-nbr** >.
- **neq** < **tcp/udp-port-nbr** > — “Not Equal”; to have a match with the ACE entry, the TCP or UDP source port number in a packet must not be equal to < **tcp/udp-port-nbr** >.
- **range** < **start-port-nbr** > < **end-port-nbr** > — For a match with the ACE entry, the TCP or UDP source-port number in a packet must be in the range < **start-port-nbr** > < **end-port-nbr** >.

Port Number or Well-Known Port Name:

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

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

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

[comparison-operator < tcp-dest-port >] [established]

[comparison-operator < udp-dest-port >]

This option, if used, is entered immediately after the < DA > entry. To specify a TCP or UDP port number, (1) select a comparison operator and (2) enter the port number or a well-known port name.

Comparison Operators and Well-Known Port Names —

These are the same as are used with the TCP/UDP source-port options, and are listed earlier in this command description.

[established] — *This option applies only where TCP is the configured IP protocol type. It blocks the synchronizing packet associated with establishing a TCP connection in one direction on a VLAN while allowing all other IP traffic for the same type of connection in the opposite direction. For example, a Telnet connect requires TCP traffic to move both ways between a host and the target device. Simply applying a Deny to inbound Telnet traffic on a VLAN would prevent Telnet sessions in either direction because responses to outbound requests would be blocked. However, by using the **established** option, inbound Telnet traffic arriving in response to outbound Telnet requests would be permitted, but inbound Telnet traffic trying to establish a connection would be denied.*

Options for ICMP Traffic in Extended ACLs. This option is useful where it is necessary to permit some types of ICMP traffic and deny other types, instead of simply permitting or denying all types of ICMP traffic. That is, an ACE designed to permit or deny ICMP traffic can optionally include an ICMP type and code value to permit or deny an individual type of ICMP packet while not addressing other ICMP traffic types in the same ACE. As an optional alternative, the ACE can include the name of an ICMP packet type. (For a summary of the extended ACL syntax options, refer to table 10-10 on page 10-60.)

Syntax: < deny | permit > icmp < SA > < DA > [icmp-type [icmp-code]
< deny | permit > icmp < SA > < DA > [icmp-type-name]

[] []

*In an extended ACL using **icmp** as the packet protocol type (see above), you can optionally specify an individual ICMP packet type or packet type/code pair to further define the criteria for a match. This option, if used, is entered immediately after the destination IP address (DA) entry. The following example shows two ACEs entered in a Named ACL context:*

```
#permit icmp any any host-unknown  
#permit icmp any any 3 7
```

[icmp-type [icmp-code]]

This option identifies an individual ICMP packet type as criteria for permitting or denying that type of ICMP traffic in an ACE.

- **icmp-type** — This value is in the range of 0 - 255 and corresponds to an ICMP packet type.
- **icmp-code** — This value is in the range of 0 - 255 and corresponds to an ICMP code for an ICMP packet type.

For more information on ICMP type names, visit the Internet Assigned Numbers Authority (IANA) website at www.iana.com, click on “Protocol Number Assignment Services”, and then go to the selections under “Internet Control Message Protocol (ICMP) Parameters”.

[*icmp-type-name*]

These name options are an alternative to the [icmp-type [icmp-code]] methodology described above. For more information, visit the IANA website cited above.

administratively-prohibited	net-tos-unreachable
alternate-address	net-unreachable
conversion-error	network-unknown
dod-host-prohibited	no-room-for-option
dod-net-prohibited	option-missing
echo	packet-too-big
echo-reply	parameter-problem
general-parameter-problem	port-unreachable
host-isolated	precedence-unreachable
host-precedence-unreachable	protocol-unreachable
host-redirect	reassembly-timeout
host-tos-redirect	redirect
host-tos-unreachable	router-advertisement
host-unknown	router-solicitation
host-unreachable	source-quench
information-reply	source-route-failed
information-request	time-exceeded
mask-reply	timestamp-reply
mask-request	timestamp-request
mobile-redirect	traceroute
net-redirect	ttl-exceeded
net-tos-redirect	unreachable

Option for IGMP in Extended ACLs. This option is useful where it is necessary to permit some types of IGMP traffic and deny other types instead of simply permitting or denying all types of IGMP traffic. That is, an ACE designed to permit or deny IGMP traffic can optionally include an IGMP packet type to permit or deny an individual type of IGMP packet while not addressing other IGMP traffic types in the same ACE. (For a summary of the extended ACL syntax options, refer to table 10-10 on page 10-60.)

Syntax: < permit | deny > igmp < SA > < DA > [igmp-type]

*In an extended ACL using **igmp** as the packet protocol type, you can optionally specify an individual IGMP packet type to further define the criteria for a match. This option, if used, is entered immediately after the destination IP addressing entry. The following example shows an IGMP ACE entered in the Named ACL context:*

```
ProCurve(config-ext-nacl)# permit igmp any  
any host-query
```

[igmp-type]

The complete list of IGMP packet-type options includes:

dvmrp	trace	mtrace-request
host-query	v2-host-report	v3-host-report
host-report	v2-host-leave	
pim	mtrace-reply	

For more information on IGMP packet types, visit the Internet Assigned Numbers Authority (IANA) website at www.iana.com, click on “Protocol Number Assignment Services”, and then go to the selections under “Internet Group Management Protocol (IGMP) Type Numbers”.

Example of a Named, Extended ACL. Suppose that you want to implement these policies on a switch configured for IP routing and membership in VLANs 10, 20, and 30:

- A. Permit Telnet traffic from 10.10.10.44 to 10.10.20.78, deny all other 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 10-18, below.)
- B. Permit FTP traffic from IP address 10.10.20.100 (on VLAN 20) to 10.10.30.55 (on VLAN 30). Deny FTP traffic from other hosts on network 10.10.20.0 to any destination, but permit all other IP traffic.




Figure 10-18. Example of an Extended ACL

```
A (Refer to figure 10-18 on page  
ProCurve(config)# ip access-list extended Extended-List-01  
| ProCurve(config-ext-nacl)# permit tcp host 10.10.10.44 host  
| 10.10.20.78 eq telnet  
| ProCurve(config-ext-nacl)# deny ip 10.10.10.1/24 10.10.20.1/24  
| ProCurve(config-ext-nacl)# permit ip any any  
| ProCurve(config-ext-nacl)# exit  
| ProCurve(config)# vlan 10 ip access-group Extended-List in  
B (Refer to figure 10-18 on page  
ProCurve(config)# ip access-list extended Extended-List-02  
| ProCurve(config-ext-nacl)# permit tcp host 10.10.20.100 host  
| 10.10.30.55 eq ftp  
| ProCurve(config-ext-nacl)# deny tcp 10.10.20.1/24 any eq ftp log  
| ProCurve(config-ext-nacl)# permit ip any any  
| ProCurve(config-ext-nacl)# exit  
| ProCurve(config)# vlan 20 ip access-group Extended-List-02 in
```

Figure 10-19. Example of Configuration Commands for Extended ACLs

Configuring Numbered, Extended ACLs

This section describes the commands for performing the following in a numbered, extended ACL:

- creating the ACL by entering the first ACE in the list
- appending a new ACE to the end of an existing ACL

For other ACL topics, refer to the following:

Topic	Page
configuring named, standard ACLs	10-53
configuring numbered, standard ACLs	10-56
configuring named, extended ACLs	10-62
applying or removing an ACL on an interface	10-81
deleting an ACL	10-85
editing an ACL	10-86
sequence numbering in ACLs	10-87
including remarks in an ACL	10-92
displaying ACL configuration data	10-96
creating or editing ACLs offline	10-104
enabling ACL “Deny” logging	10-109

Creating or Adding to an Extended, Numbered ACL. This command is an alternative to using **ip access-list extended < name-str >** and does not use the Named ACL (**nacl**) context. (For an extended ACL syntax summary, refer to table 10-10 on page 10-60.)

Syntax: access-list < 100-199 > < deny | permit > < ip | ip-protocol | ip-protocol-nbr >
< any | host < SA > | SA/mask-length | SA < mask >>
< any | host < DA > | DA/mask-length | DA < mask >>
[precedence < 0 - 7 | precedence-name >]
[tos < tos-bit-setting >]
[log]

*If the ACL does not already exist, this command creates the specified ACL and its first ACE. If the ACL already exists, the new ACE is appended to the end of the configured list of explicit ACEs. In the default configuration, the ACEs in an ACL will automatically be assigned consecutive sequence numbers in increments of 10 and can be renumbered with **resequence** (page 10-91).*

Note: To insert a new ACE between two existing ACEs in an extended, numbered ACL:

- a. Use **ip access list extended < 100 - 199 >** to open the ACL as a named ACL.
- b. Enter the desired sequence number along with the ACE statement you want.

(Refer to the “Numbered ACLs” list item on page 10-49.)

For a match to occur, a packet must have the source and destination IP addressing criteria specified in the ACE, as well as:

- *the protocol-specific criteria configured in the ACE, including any included, optional elements (described later in this section)*
- *any (optional) precedence and/or ToS settings configured in the ACE*

< 100-199 >

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

< deny | permit >

*Specifies whether to deny (**drop**) or permit (forward) a packet that matches the criteria specified in the ACE, as described below.*

< ip | ip-protocol | ip-protocol-nbr >

Specifies the packet protocol type required for a match. An extended ACL must include one of the following:

- **ip** — any IP packet.
 - **ip-protocol** — any one of the following IP protocol names:

ip-in-ip	ipv6-in-ip	gre	esp	ah
ospf	pim	vrrp	sctp	tcp*
udp*	icmp*	igmp*		
 - **ip-protocol-nbr** — the IPv4 IP protocol number of an IP packet type, such as “8” for Exterior Gateway Protocol or 121 for Simple Message Protocol. (For a listing of IP protocol numbers and their corresponding protocol names, refer to the IANA “Protocol Number Assignment Services” at www.iana.com.) (Range: 0 - 255)
- * For TCP, UDP, ICMP, and IGMP, additional criteria can be specified, as described later in this section.

< any | host < SA > | SA/mask-length | SA < mask >>

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 < SA >** — 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.
- **SA/mask-length** or **SA < mask >** — Specifies packets received from an SA, where the SA is 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 10-50.

SA Mask Application: *The mask is applied to the SA in the ACL to define which bits in a packet's source SA must exactly match the IP address configured in the ACL and which bits need not match.*

Example: *10.10.10.1/24 and 10.10.10.1 0.0.0.255 both define any IP address in the range of 10.10.10.(1-255).*

Note: *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 10-36.*

< any | host < DA > | DA/mask-length >

This is the second instance of IP addressing in an extended ACE. It follows the first (SA) instance, described earlier, and 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 < SA >.

- **any** — *Allows routed IP packets to any DA.*
- **host < DA >** — *Specifies only packets having DA as the destination IP address. Use this criterion when you want to match only the IP packets for a single DA.*
- **DA/mask-length** or **DA < mask >** — *Specifies packets intended for a destination address, where the address is either a subnet or a group of IP addresses. The mask format can be in either dotted-decimal format or CIDR format (number of significant bits). Refer to "Using CIDR Notation To Enter the ACL Mask" on page 10-50.*

DA Mask Application: *The mask is applied to the DA in the ACL to define which bits in a packet's DA must exactly match the DA configured in the ACL and which bits need not match. See also the above example and note.*

[precedence < 0 - 7 | precedence-name >]

This option causes the ACE to match packets with the specified IP precedence value. Values can be entered as the following IP precedence numbers or alphanumeric names:

0	or	routine
1	“	priority
2	“	immediate
3	“	flash
4	“	flash-override
5	“	critical
6	“	internet (for internetwork control)
7	“	network (for network control)

Note: *The precedence criteria described in this section are applied in addition to any other selection criteria configured in the same ACE.*

[tos]

This option can be used after the DA to cause the ACE to match packets with the specified IP Type-of-Service (ToS) setting. ToS values can be entered as the following numeric settings or, in the case of 0, 2, 4, and 8, as alphanumeric names:

0	or	normal
2	“	max-reliability
4	“	max-throughput
6		
8	“	minimize-delay
10		
12		
14		

Note: *The ToS criteria in this section are applied in addition to any other criteria configured in the same ACE.*

[log]

Optional; generates an Event 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 10-111)*

Additional Options for TCP and UDP Traffic. An ACE designed to permit or deny TCP or UDP traffic can optionally include port number criteria for either the source or destination, or both. Use of TCP criteria also allows the **established** option for controlling TCP connection traffic. (For a summary of the extended ACL syntax options, refer to table 10-10 on page 10-60.)

Syntax: access-list < 100 - 199 > < deny | permit > < tcp | udp >
 < SA > [comparison-operator < tcp/udp-src-port >]

 < DA > [comparison-operator < tcp-dest-port >] [established]
 < DA > [comparison-operator < udp-dest-port >]

This source-port and destination-port TCP/UDP criteria is identical to the criteria described for TCP/UDP use in named, extended ACLs, beginning on page 10-68.

Additional Options for ICMP Traffic. This option is useful where it is necessary to permit some types of ICMP traffic and deny other types, instead of simply permitting or denying all types of ICMP traffic. That is, an ACE designed to permit or deny ICMP traffic can optionally include an ICMP type and code value to permit or deny an individual type of ICMP packet while not addressing other ICMP traffic types in the same ACE. As an optional alternative, the ACE can include the name of an ICMP packet type. (For a summary of the extended ACL syntax options, refer to table 10-10 on page 10-60.)

Syntax: access-list < 100 - 199 > < deny | permit > icmp < SA > < DA >
 [[icmp-type [icmp-code]] | [icmp-type-name]]

The ICMP “type” and “code” criteria are identical to the criteria described for ICMP in named, extended ACLs, beginning on page 10-70.

Additional Option for IGMP. This option is useful where it is necessary to permit some types of IGMP traffic and deny other types, instead of simply permitting or denying all types of IGMP traffic. That is, an ACE designed to permit or deny IGMP traffic can optionally include an IGMP packet type to permit or deny an individual type of IGMP packet while not addressing other IGMP traffic types in the same ACE. (For a summary of the extended ACL syntax options, refer to table 10-10 on page 10-60.)

Syntax: access-list < 100 - 199 >
< deny | permit > igmp < src-ip > < dest-ip > [igmp-type]

The IGMP “type” criteria is identical to the criteria described for IGMP in named, extended ACLs, beginning on page 10-72.

Adding or Removing an ACL Assignment On an Interface

Filtering Routed IP Traffic

For a given VLAN interface on a switch configured for routing, you can assign an ACL as a RACL to filter inbound IP traffic and another ACL as a RACL to filter outbound IP traffic. You can also assign one ACL for both inbound and outbound RACLs, and for assignment to multiple VLANs. For limits and operating rules, refer to “ACL Configuration and Operating Rules” on page 10-33.

Syntax: [no] vlan < vid > ip access-group < identifier > < in | out >
where: < identifier > = either a ACL name or an ACL ID number.

Assigns an ACL to a VLAN as an RACL to filter routed IP traffic entering or leaving the switch on that VLAN. You can use either the global configuration level or the VLAN context level to assign or remove an RACL.

Note: *The switch allows you to assign a nonexistent ACL name or number to a VLAN. In this case, if you subsequently configure an ACL with that name or number, it automatically becomes active on the assigned VLAN. Also, if you delete an assigned ACL from the switch without subsequently using the “no” form of this command to remove the assignment to a VLAN, the ACL assignment remains and will automatically activate any new ACL you create with the same identifier (name or number).*

Access Control Lists (ACLs)

Adding or Removing an ACL Assignment On an Interface

ProCurve(config)# vlan 20 ip access-group My-List in	←	Enables an RACL from the Global Configuration Level
ProCurve(config)# vlan 20		
ProCurve(vlan-20)# ip access-group 155 out	←	Enables an RACL from a VLAN Context.
ProCurve(vlan-20)# exit		
ProCurve(config)# no vlan 20 ip access-group My-List in	←	Disables an RACL from the Global Configuration Level
ProCurve(config)# vlan 20		
ProCurve(vlan-20)# no ip access-group 155 out	←	Disabling an RACL from a VLAN Context.
ProCurve(vlan-20)# exit		

Figure 10-20. Methods for Enabling and Disabling RACLs

Filtering IP Traffic Inbound on a VLAN

For a given VLAN interface, you can assign an ACL as a VACL to filter any IP traffic entering the switch on that VLAN. You can also use the same ACL for assignment to multiple VLANs. For limits and operating rules, refer to “ACL Configuration and Operating Rules” on page 10-33.

Syntax: [no] vlan < vid > ip access-group < identifier > vlan
where: < identifier > = either a ACL name or an ACL ID number.

Assigns an ACL as a VACL to a VLAN to filter any IP traffic entering the switch on that VLAN. You can use either the global configuration level or the VLAN context level to assign or remove a VACL.

Note: *The switch allows you to assign a nonexistent ACL name or number to a VLAN. In this case, if you subsequently configure an ACL with that name or number, it automatically becomes active on the assigned VLAN. Also, if you delete an assigned ACL from the switch without subsequently using the “no” form of this command to remove the assignment to a VLAN, the ACL assignment remains and will automatically activate any new ACL you create with the same identifier (name or number).*

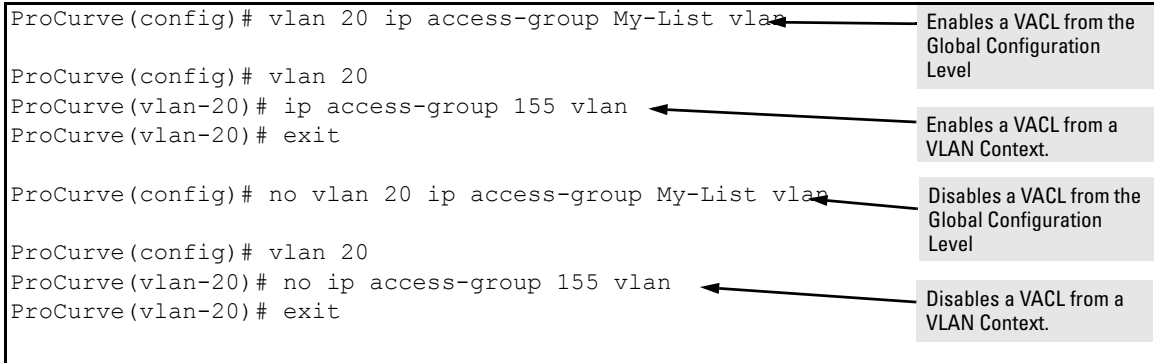


Figure 10-21. Methods for Enabling and Disabling VACLs

Filtering Inbound IP Traffic Per Port

For a given port, port list, or static port trunk, you can assign an ACL as a static port ACL to filter any IP traffic entering the switch on that interface. You can also use the same ACL for assignment to multiple interfaces. For limits and operating rules, refer to “ACL Configuration and Operating Rules” on page 10-33.

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

Assigns an ACL as a static port ACL to a port, port list, or static trunk to filter any IP traffic entering the switch on that interface. You can use either the global configuration level or the interface context level to assign or remove a static port ACL.

Note: *The switch allows you to assign a nonexistent ACL name or number to an interface. In this case, if you subsequently configure an ACL with that name or number, it automatically becomes active on the assigned interface. Also, if you delete an assigned ACL from the switch without subsequently using the “no” form of this command to remove the assignment to an interface, the ACL assignment remains and will automatically activate any new ACL you create with the same identifier (name or number).*

ProCurve(config)# interface b10 ip access-group My-List in	←	Enables a static port ACL from the Global Configuration level.
ProCurve(config)# interface b10		
ProCurve(eth-b10)# ip access-group 155 in	←	Enables a static port ACL from a port context.
ProCurve(eth-b10)# exit		
ProCurve(config)# no interface b10 ip access-group My-List in	←	Disables a static port ACL from the Global Configuration level.
ProCurve(config)# interface b10		
ProCurve(eth-b10)# no ip access-group 155 in	←	Uses a VLAN context to disable a static port ACL.
ProCurve(eth-b10)# exit		

Figure 10-22. Methods for Enabling and Disabling ACLs

Deleting an ACL

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

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

no access-list < 1 - 99 | 100 - 199 >

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

Note: *Deleting an ACL does not delete any assignment of that ACL's identifier on a specific interface. Creating a new ACL using an identifier that is already configured on an interface causes the switch to automatically activate that ACL. If you need to remove an ACL identifier assignment on an interface, refer to "Adding or Removing an ACL Assignment On an Interface" on page 10-81*

Editing an Existing ACL

The CLI provides the capability for editing in the switch by using sequence numbers to insert or delete individual ACEs. An offline method is also available. This section describes using the CLI for editing ACLs. To use the offline method for editing ACLs, refer to “Creating or Editing ACLs Offline” on page 10-104.

Using the CLI To Edit ACLs

You can use the CLI to delete individual ACEs from anywhere in an ACL, append new ACEs to the end of an ACL, and insert new ACEs anywhere within an ACL.

General Editing Rules

- **Named ACLs:**
 - When you enter a new ACE in a named ACL without specifying a sequence number, the switch inserts the ACE as the last entry in the ACL.
 - When you enter a new ACE in a named ACL and include a sequence number, the switch inserts the ACE according to the position of the sequence number in the current list of ACEs.
- **Numbered ACLs:** When using the **access-list < 1 - 99 | 100 - 199 >** command to create or add ACEs to a numbered ACL, each new ACE you enter is added to the end of the current list. (This command does not offer a **< seq-# >** option for including a sequence number to enable inserting an ACE at other points in the list.) Note, however, that once a numbered list has been created, you have the option of accessing it in the same way as a named list by using the **ip access-list < standard | extended >** command. This enables you to edit a numbered list in the same way that you would edit a named list. (See the next item in this list.)
- You can delete any ACE from any ACL (named or numbered) by using the **ip access-list** command to enter the ACLs context, and then using the **no < seq-# >** command (page 10-90).

- Deleting the last ACE from an ACL leaves the ACL in memory. In this case, the ACL is “empty” and cannot perform any filtering tasks. (In any ACL the Implicit Deny does not apply unless the ACL includes at least one explicit ACE.)

Sequence Numbering in ACLs

The ACEs in any ACL are sequentially numbered. In the default state, the sequence number of the first ACE in a list is “10” and subsequent ACEs are numbered in increments of 10. For example, the following **show run** output lists three ACEs with default numbering in a list named “My-List”:

```
ip access-list standard "My-List"  
 10 permit 10.10.10.25 0.0.0.0  
 20 permit 10.20.10.117 0.0.0.0  
 30 deny 10.20.10.1 0.0.0.255  
exit
```

Figure 10-23. Example of the Default Sequential Numbering for ACEs

You can add an ACE to the end of a named or numbered ACL by using either **access-list** for numbered ACLs or **ip access-list** for named ACLs:

```
ProCurve(config)# access-list 2 permit any  
  
ProCurve(Config)# ip access-list standard My-list  
ProCurve(Config-ext-nacl)# permit ip any host 10.10.10.125
```

← Appends an ACE to the end of a standard, numbered ACL.

↗ Enters the context of an extended ACL and appends an ACE to the end of the list.

Figure 10-24. Examples of Adding an ACE to the end of Numbered or Named ACLs

For example, to append a fourth ACE to the end of the ACL in figure 10-23:

```
ProCurve(config)# ip access-list standard My-List
ProCurve(config-std-nacl)# permit any
ProCurve(config-std-nacl)# show run
.
.
.
ip access-list standard "My-List"
 10 permit 10.10.10.25 0.0.0.0
 20 permit 10.20.10.117 0.0.0.0
 30 deny 10.20.10.1 0.0.0.255
 40 permit 0.0.0.0 255.255.255.255
exit
```

Figure 10-25. Example of Appending an ACE to an Existing List

Note

When using the **access-list < 1 - 99 | 100 - 199 > < permit | deny > < SA >** command to create an ACE for a numbered ACL, the ACE is always added to the end of the current list and given the appropriate sequence number. However, once a numbered list has been created, you can use the **ip access-list** command to open it as a named ACL and specify a nondefault sequence number, as described in the next section.

Inserting an ACE in an Existing ACL

This action uses a sequence number to specify where to insert a new ACE into an existing sequence of ACLs.

Syntax: ip access-list < standard | extended > < name-str | 1 - 99 | 100 - 199 >

```
<1-2147483647> < permit | deny > < standard-acl-ip-criteria > [ log ]
<1-2147483647> < permit | deny > < extended-acl-ip-criteria > [ options ]
```

The first command enters the “Named-ACL” context for the specified ACL. The remaining two commands insert a new ACE in a standard or extended ACL, respectively. (For details on these criteria and options, refer to table 10-1, “Command Summary for Standard ACLs” —page 10-6, and table 10-2, “Command Summary for Extended ACLs” —page 10-8.)

To insert a new ACE between existing ACEs in a list:

1. Use **ip access-list** to enter the “Named-ACL” (**nacl**) context of the ACE. *This applies regardless of whether the ACE was originally created as a numbered ACL or a named ACL.*

2. Begin the ACE command with a sequence number that identifies the position you want the ACE to occupy. (The sequence number range is 1-2147483647.)
3. Complete the ACE with the command syntax appropriate for the type of ACL you are editing.

For example, inserting a new ACE between the ACEs numbered 10 and 20 in figure 10-25 requires a sequence number in the range of 11-19 for the new ACE.

```
ProCurve(config)# ip access-list standard My-List
ProCurve(config-std-nacl)# 15 deny 10.10.10.1/24
ProCurve(config-std-nacl)# show run
.
.
.
ip access-list standard "My-List"
 10 permit 10.10.10.25 0.0.0.0
 15 deny 10.10.10.1 0.0.0.255
 20 permit 10.20.10.117 0.0.0.0
 30 deny 10.20.10.1 0.0.0.255
 40 permit 0.0.0.0 255.255.255.255
exit
```

Figure 10-26. Example of Inserting an ACE in an Existing ACL

In the following example, the first two ACEs entered become lines 10 and 20 in the list. The third ACE entered is configured with a sequence number of 15 and is inserted between lines 10 and 20.

```
ProCurve(config)# ip access-list standard List-01
ProCurve(config-std-nacl)# permit 10.10.10.1/24
ProCurve(config-std-nacl)# deny 10.10.1.1/16
ProCurve(config-std-nacl)# 15 permit 10.10.20.1/24
ProCurve(config-std-nacl)# show run

Running configuration:
. . .
ip access-list standard "List-01"
 10 permit 10.10.10.1 0.0.0.255
 15 permit 10.10.20.1 0.0.0.255
 20 deny 10.10.1.1 0.0.255.255
exit
```

Figure 10-27. Example of Inserting an ACE into an Existing Sequence

Deleting an ACE from an Existing ACL

This action uses ACL sequence numbers to delete ACEs from an ACL.

Syntax: ip access-list < standard | extended > < name-str | 1 - 99 | 100 - 199 >
no < seq-# >

The first command enters the “Named-ACL” context for the specified ACL. The no command deletes the ACE corresponding to the sequence number entered. (Range: 1 - 2147483647)

1. To find the sequence number of the ACE you want to delete, use **show run** or **show access-list < name-str | 1 - 99 | 100-199 >** to view the ACL.
2. Use **ip access-list** to enter the “Named-ACL” (**nacl**) context of the ACE. *This applies regardless of whether the ACE was originally created as a numbered ACL or a named ACL.*
3. In the “Named-ACL” context, type **no** and enter the sequence number of the ACE you want to delete.

Figure 10-28 illustrates the process for deleting an ACE from a list:

```
ProCurve(config)# show run
. . .
ACL Before Deleting an ACE
ip access-list standard "My-List"
 10 permit 10.10.10.25 0.0.0.0
 15 deny 10.10.10.1 0.0.0.255
 20 permit 10.20.10.117 0.0.0.0
 30 deny 10.20.10.1 0.0.0.255
 40 permit 0.0.0.0 255.255.255.255
  exit
ProCurve(config)# ip access-list standard My-List
ProCurve(config-std-nacl)# no 20
ProCurve(config-std-nacl)# show run
. . .
ACL After Deleting the ACE at Line 20
ip access-list standard "My-List"
 10 permit 10.10.10.25 0.0.0.0
 15 deny 10.10.10.1 0.0.0.255
 30 deny 10.20.10.1 0.0.0.255
 40 permit 0.0.0.0 255.255.255.255
  exit
```

This command enters the “Named-ACL” (nacl) context for “My-List”.

This command deletes the ACE at line 20.

The ACE at line 20 has been removed.

Figure 10-28. Example of Deleting an ACE from Any ACL

Resequencing the ACEs in an ACL

This action reconfigures the starting sequence number for ACEs in an ACL, and resets the numeric interval between sequence numbers for ACEs configured in the ACL.

Syntax: `ip access-list resequence < name-str | 1 - 99 | 100 - 199 >
< starting-seq-# > < interval >`

Resets the sequence numbers for all ACEs in the ACL.

< starting-seq-# > : Specifies the sequence number for the first ACE in the list. (Default: 10; Range: 1 - 2147483647)

< interval > : Specifies the interval between sequence numbers for the ACEs in the list. (Default: 10; Range: 1 - 2147483647)

1. To view the current sequence numbering in an ACE, use **show run** or **show access-list < name-str | 1 - 99 | 100-199 >**.
2. Use the command syntax (above) to change the sequence numbering.

This example resequences the “My-List” ACL at the bottom of figure 10-28 so that the list begins with line 100 and uses a sequence interval of 100.

```
ProCurve(config)# show run
. . .
ip access-list standard "My-List"
 10 permit 10.10.10.25 0.0.0.0
 15 deny 10.10.10.1 0.0.0.255
 30 deny 10.20.10.1 0.0.0.255
 40 permit 0.0.0.0 255.255.255.255
exit
. . .
ProCurve(config)# ip access-list resequence My-List 100 100
ProCurve(config)# show run
. . .
ip access-list standard "My-List"
 100 permit 10.10.10.25 0.0.0.0
 200 deny 10.10.10.1 0.0.0.255
 300 deny 10.20.10.1 0.0.0.255
 400 permit 0.0.0.0 255.255.255.255
exit
```

Figure 10-29. Example of Viewing and Resequencing an ACL

Attaching a Remark to an ACE

A remark is numbered in the same way as an ACE, and uses the same sequence number as the ACE to which it refers. This operation requires that the remark for a given ACE be entered prior to entering the ACE itself.

Syntax: access-list < 1 - 99 | 100 - 199 > remark < remark-str >

This syntax appends a remark to the end of a numbered ACL and automatically assigns a sequence number to the remark. The next command entry should be the ACE to which the remark belongs. (The new ACE will automatically be numbered with the same sequence number as was used for the preceding remark.)

Syntax: ip access-list < standard | extended > < name-str | 1-99 | 100-199 >
[seq-#] remark < remark-str >
no < seq-# > remark

*This syntax applies to both named and numbered ACLs. Without an optional sequence number, the remark is appended to the end of the list and automatically assigned a sequence number. When entered with an optional sequence number, the remark is inserted in the list according to the numeric precedence of the sequence number. The **no** form of the command deletes the indicated remark, but does not affect the related ACE.*

To associate a remark with a specific ACE, enter the remark first, and then enter the ACE.

- Entering a remark without a sequence number and then entering an ACE without a sequence number results in the two entries being automatically paired with the same sequence number and appended to the end of the current ACL.*
- Entering a remark with a sequence number and then entering an ACE with the same sequence number results in the two entries being paired together and positioned in the list according to the sequence number they share.*

Note

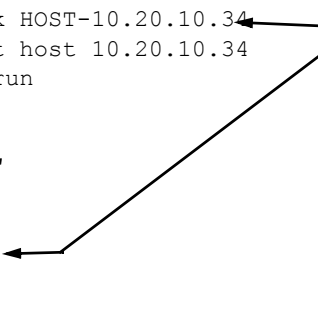
After a numbered ACL has been created (using **access-list < 1 - 99 | 100 - 199 >**), it can be managed as either a named or numbered ACL. For example, in an existing ACL with a numeric identifier of “115”, either of the following command sets adds an ACE denying IP traffic from any IP source to a host at 10.10.10.100:

```
ProCurve(config)# access-list 115 deny ip host
10.10.10.100

ProCurve(config)# ip access-list extended 115
ProCurve(config-ext-nacl)# deny ip any 10.10.10.100
```

Appending Remarks and Related ACEs to the End of an ACL. To include a remark for an ACE that will be appended to the end of the current ACL, enter the remark first, then enter the related ACE. This results in the remark and the subsequent ACE having the same sequence number. For example, to add remarks using the “Named-ACL” (**nacl**) context:

```
ProCurve(config)# ip access-list standard My-List
ProCurve(config-std-nacl)# permit host 10.10.10.15
ProCurve(config-std-nacl)# deny 10.10.10.1/24
ProCurve(config-std-nacl)# remark HOST-10.20.10.34
ProCurve(config-std-nacl)# permit host 10.20.10.34
ProCurve(config-std-nacl)# show run
. . .
hostname "ProCurve"
ip access-list standard "My-List"
 10 permit 10.10.10.15 0.0.0.0
 20 deny 10.10.10.1 0.0.0.255
 30 remark "HOST-10.20.10.34"
 30 permit 10.20.10.34 0.0.0.0
exit
```



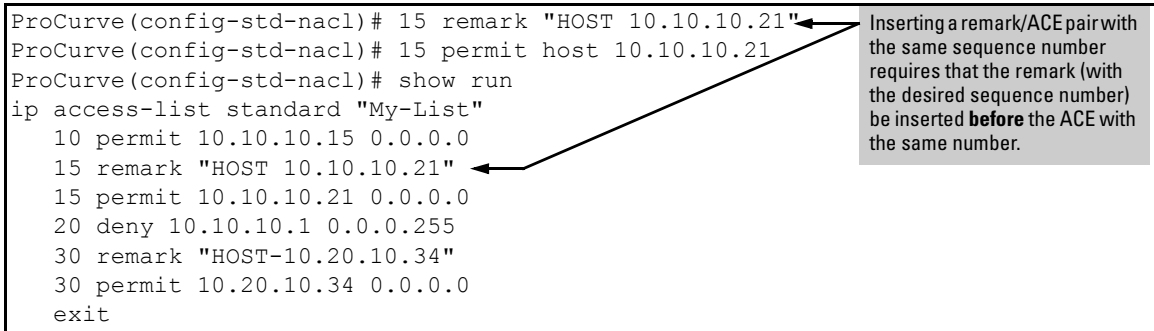
The remark is assigned the same number that the immediately following ACE (“30” in this example) is assigned when it is automatically appended to the end of the list. This operation applies where new remarks and ACEs are appended to the end of the ACL and are automatically assigned a sequence number.

Figure 10-30. Example of Appending a Remark and Its Related ACE to the End of an ACL

(You can also perform the operation illustrated in figure 10-30 by using the numbered, **access-list < 1 - 99 | 100 - 199 >** syntax shown at the beginning of this section.)

Inserting Remarks and Related ACEs Within an Existing List. To insert an ACE with a remark within an ACL by specifying a sequence number, insert the numbered remark first, then, using the same sequence number, insert the ACE. (This operation applies only to ACLs accessed using the “Named-ACL” (**nacl**) context.) For example:

```
ProCurve(config-std-nacl)# 15 remark "HOST 10.10.10.21"
ProCurve(config-std-nacl)# 15 permit host 10.10.10.21
ProCurve(config-std-nacl)# show run
ip access-list standard "My-List"
 10 permit 10.10.10.15 0.0.0.0
 15 remark "HOST 10.10.10.21"
 15 permit 10.10.10.21 0.0.0.0
 20 deny 10.10.10.1 0.0.0.255
 30 remark "HOST-10.20.10.34"
 30 permit 10.20.10.34 0.0.0.0
exit
```



Inserting a remark/ACE pair with the same sequence number requires that the remark (with the desired sequence number) be inserted **before** the ACE with the same number.

Inserting a Remark for an ACE that Already Exists in an ACL. If a sequence number is already assigned to an ACE in a list, you cannot insert a remark by assigning it to the same number. (To configure a remark with the same number as a given ACE, the remark must be configured first.) To assign a remark to the same number as an existing ACE:

1. Delete the ACE.
2. Configure the remark with the number you want assigned to the pair.
3. Re-Enter the deleted ACE with the number used to enter the remark.

Removing a Remark from an Existing ACE. If you want to remove a remark, but want to retain the ACE, do the following:

1. Use the Named ACL context to enter the ACL.
2. Using **show run** or **show access-list <list-name >**, note the sequence number and content of the ACE having a remark you want to remove.
3. Delete the ACE.
4. Using the same sequence number, re-enter the ACE.

Operating Notes for Remarks

- The **resequence** command ignores “orphan” remarks that do not have an ACE counterpart with the same sequence number. For example, if:
 - a remark numbered “55” exists in an ACE
 - there is no ACE numbered “55” in the same ACL
 - **resequence** is executed on an ACL

then the remark retains “55” as its sequence number and will be placed in the renumbered version of the ACL according to that sequence number.

- Entering an unnumbered remark followed by a numbered ACE, or the reverse, creates an “orphan” remark. The unnumbered entry will be assigned a sequence number that is an increment from the last ACE in the list. The numbered entry will then be placed sequentially in the list according to the sequence number used.
- Configuring two remarks without either sequence numbers or an intervening, unnumbered ACE results in the second remark overwriting the first.

```
ProCurve(config)# ip access-list standard Accounting
ProCurve(config-std-nacl)# permit host 10.10.10.115
ProCurve(config-std-nacl)# deny 10.10.10.1/24
ProCurve(config-std-nacl)# remark Marketing
ProCurve(config-std-nacl)# remark Channel_Mktg
ProCurve(config-std-nacl)# show run
.
.
.
ip access-list standard "Accounting"
 10 permit 10.10.10.115 0.0.0.0
 20 deny 10.10.10.1 0.0.0.255
 30 remark "Channel_Mktg"
exit
```

Where multiple remarks are sequentially entered for automatic inclusion at the end of an ACL, each successive remark replaces the previous one until an ACE is configured for automatic inclusion at the end of the list.

Figure 10-31. Example of Overwriting One Remark with Another

Displaying ACL Configuration Data

ACL Commands	Function	Page
show access-list	Displays a brief listing of all ACLs on the switch.	10-97
show access-list config	Display the type, identifier, and content of all ACLs configured in the switch.	10-98
show access-list vlan < vid >	List the name and type for each ACL application assigned to a particular VLAN on the switch.	10-99
show access-list ports < all port-list >	Lists the ACL static port assignment for either all ports and trunks, or for the specified ports and/or trunks.	
show access-list < acl-id >	Display detailed content information for a specific ACL.	10-10 1
show access-list resources	Displays the currently available per-slot resource availability. Refer to the appendix titled "Monitoring Resources" in the current <i>Management and Configuration Guide</i> for your switch.	
show access-list radius < all port-list >	Lists the RADIUS ACL(s) currently assigned for either all ports and trunks, or for the specified ports and/or trunks. For more on this topic, refer to chapter 7, "Configuring RADIUS Server Support for Switch Services".	
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 VLANs.

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
----  -
std   yes   List-01-Inbound
ext   no    List-02-Outbound
std   yes   55
```

In this switch, the ACL named "List-02-Outbound" exists in the configuration but is not applied to any VLANs and thus does not affect traffic.

Figure 10-32. 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 a VLAN (yes/no).
Name	Shows the identifier (name or number) assigned to each ACL configured in the switch.

Display the Content of All ACLs on the Switch

This command lists the configuration details for every ACL in the running-config file, regardless of whether any are actually assigned to filter IP traffic on specific VLANs.

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 “Creating or Editing ACLs Offline” on page 10-104.

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 "List-43"
 10 deny 10.28.236.77 0.0.0.0
 20 deny 10.29.140.107 0.0.0.0
 30 permit 0.0.0.0 255.255.255.255
 exit
ip access-list extended "111"
 10 permit tcp 10.30.133.27 0.0.0.0 0.0.0.0 255.255.255.255
 20 permit tcp 10.30.155.101 0.0.0.0 0.0.0.0 255.255.255.255
 30 deny ip 10.30.133.1 0.0.0.0 0.0.0.0 255.255.255.255 log
 40 deny ip 10.30.155.1 0.0.0.255 0.0.0.0 255.255.255.255
 exit
```

Figure 10-33. Example of an ACL Configured Syntax Listing

Display the RACL and VACL Assignments for a VLAN

This command briefly lists the identification and type(s) of RACLs and VACLs currently assigned to a particular VLAN in the running-config file. (The switch allows one inbound and one outbound RACL assignment per VLAN, plus one VACL assignment.)

Syntax: show access-list vlan < vid >

Lists any RACL and/or VACL assignments to a VLAN in the running config file.

Note

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

For example, if you assigned an extended ACL with an ACL-ID of “List-43” to filter routed IP traffic exiting from the switch on VLAN 10 and a standard VACL with an ACL-ID of “List-12” to filter all IP traffic entering the switch on VLAN 10, you could verify these assignments as shown in figure 10-34:

```
ProCurve(config)# show access-list vlan 10

Access Lists for VLAN 10

  Inbound Access List: None
  Outbound Access List: List-43
  Type: Extended
  Vlan Access List : List-12
  Type: Standard
  Connection Rate Filter Access List: None
```

Indicates that:

- There is no ACL assignment to filter routed IP traffic entering the switch on VLAN 10.
- An extended ACL with the ID of “List-43” is assigned to filter routed IP traffic exiting the switch on VLAN 10.
- A standard ACL with the ID of “List-12” is assigned to filter all IP traffic entering the switch on VLAN 10.

Applies to Connection Rate Filter ACLs. (Refer to chapter 3, Virus Throttling”).

Figure 10-34. Example of Listing the ACL Assignments for a VLAN

Display Static Port ACL Assignments

This command briefly lists the identification and type(s) of current static port ACL assignments to individual switch ports and trunks, as configured in the running-config file. (The switch allows one static port ACL assignment per port.)

Syntax: `show access-list ports <all | interface >`

Lists the current static port ACL assignments for ports and trunks in the running config file.

Note

This information also appears in the **show running** display. If you execute **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 “Port-10” to filter inbound IP traffic on switch ports B10-B11 and trunk trk1, you could verify these assignments as shown in figure 10-35.

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

Access Lists for Port B10

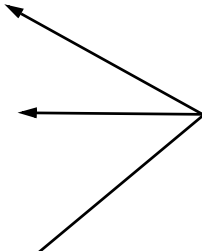
  Inbound  : 15
  Type     : Standard

Access Lists for Port B11

  Inbound  : 15
  Type     : Standard

Access Lists for Port Trk1

  Inbound  : 15
  Type     : Standard
```



Indicates that a standard ACL with the ID of “15” is assigned to filter traffic entering the switch on ports B10 and B11, and on trunk Trk1.

Figure 10-35. Example of Listing the ACL Assignments for Ports and Trunks

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 execute **write memory** after configuring an ACL, it also appears in the **show config** display.

Syntax: show access-list < acl-id >

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	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:

```
ProCurve(config)# show access-list 1

Access Control Lists

Name: 1
Type: Standard
Applied: Yes

SEQ  Entry
-----
10   Action: deny (log)
     IP      : 10.28.236.77      Mask: 0.0.0.0

20   Action: deny
     IP      : 10.29.140.107    Mask: 0.0.0.0

30   Action: permit
     IP      : 0.0.0.0          Mask: 255.255.255.255
```

Indicates whether the ACL is applied to an interface.

Figure 10-36. Example of a Listing a Standard ACL

Access Control Lists (ACLs) Displaying ACL Configuration Data

```
ProCurve(config)# show access-list List-120

Access Control Lists

Name: List-120
Type: Extended
Applied: No

SEQ  Entry
-----
10   Action: permit
     Remark: Telnet Allowed
     Src IP: 10.30.133.27      Mask: 0.0.0.0      Port (s): eq 23
     Dst IP: 0.0.0.0         Mask: 255.255.255.255  Port (s):
     Proto : TCP (Established)
     TOS   : -              Precedence: routine

20   Action: deny (log)
     Src IP: 10.30.133.1      Mask: 0.0.0.255    Port (s):
     Dst IP: 0.0.0.0         Mask: 255.255.255.255  Port (s):
     Proto : IP
     TOS   : -              Precedence: -

30   Action: permit
     Src IP: 0.0.0.0          Mask: 255.255.255.255  Port (s):
     Dst IP: 0.0.0.0         Mask: 255.255.255.255  Port (s):
     Proto : IP
     TOS   : -              Precedence: -
```

Indicates whether the ACL is applied to an interface.

Indicates source and destination entries in the ACL.

Empty field indicates that the destination TCP port can be any value.

Figure 10-37. Examples of Listings Showing the Content of Standard and Extended ACLs

Table 10-11. Descriptions of Data Types Included in Show Access-List < acl-id > Output

Field	Description
Name	The ACL identifier. Can be a number from 1 to 199, or a name.
Type	Standard or Extended. The former uses only source 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 a port or VLAN interface. “No” means the ACL exists in the switch configuration, but has not been applied to any interface, and is therefore not in use.
SEQ	The sequential number of the Access Control Entry (ACE) in the specified ACL.
Entry	Lists the content of the ACEs in the selected ACL.
Action	Permit (forward) or deny (drop) a packet when it is compared to the criteria in the applicable ACE and found to match. Includes the optional log option, if used, in deny actions.
Remark	Displays any optional remark text configured for the selected ACE.
IP	Used for Standard ACLs: The source IP address to which the configured mask is applied to determine whether there is a match with a packet.
Src IP	Used for Extended ACLs: Same as above.
Dst IP	Used for Extended ACLs: The source and destination 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. For TCP protocol selections, includes the established option, if configured.
Port(s)	Used only in extended ACLs to show any TCP or UDP operator and port number(s) included in the ACE.
TOS	Used only in extended ACLs to indicate Type-of-Service setting, if any.
Precedence	Used only in extended ACLs to indicate the IP precedence setting, if any.

Display All ACLs and Their Assignments in the Routing Switch Startup-Config File and Running-Config File

The **show config** and **show running** commands include in their listings any configured ACLs and any ACL assignments to VLANs. Refer to figure 10-12 (page 10-46) for an example. Remember that **show config** lists the startup-config file and **show running** lists the running-config file.

Creating or Editing ACLs Offline

The section titled “Editing an Existing ACL” on page 10-86 describes how to use the CLI to edit an ACL, and is most applicable in cases where the ACL is short or there is only a minor editing task to perform. The offline method provides a useful alternative to using the CLI for creating or extensively editing a large ACL. This section describes how to:

- move an existing ACL to a TFTP server
- use a text (.txt) file format to create a new ACL or edit an existing ACL offline
- use TFTP to load an offline ACL into the switch’s running-config

For longer ACLs that may be difficult or time-consuming to accurately create or edit in the CLI, you can use the offline method described in this section.

Note

Beginning with software release K_12_XX or later, **copy** commands that used either **tftp** or **xmodem**, also include an option to use **usb** as a source or destination device for file transfers. So although the following example highlights **tftp**, bear in mind that **xmodem** or **usb** can also be used to transfer ACLs to and from the switch.

Creating or Editing an ACL Offline

The Offline Process

1. Begin by doing one of the following:
 - To edit one or more existing ACLs, use **copy command-output tftp** to copy the current version of the ACL configuration to a file in your TFTP server. For example, to copy the ACL configuration to a file named **acl-02.txt** in the TFTP directory on a server at 10.28.227.2:

```
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 (.txt) file in the appropriate directory on a TFTP server accessible to the switch.
2. Use a text editor to create or edit the ACL(s) in the ***.txt** ASCII file format.

If you are replacing an ACL on the switch with a new ACL that uses the same number or name syntax, begin the command file with a **no ip access-list** command to remove the earlier version of the ACL from the switch's running-config file. Otherwise, the switch will append the new ACEs in the ACL you download to the existing ACL. For example, if you planned to use the **copy** command to *replace* ACL "List-120", you would place this command at the beginning of the edited file:

```
no ip access-list extended List-120
```

<pre>no ip access-list extended List-120 ip access-list extended "List-120" 10 remark "THIS ACE ALLOWS TELNET" 10 permit tcp 10.30.133.27 0.0.0.0 eq 23 0.0.0.0 255.255.255. 20 deny ip 10.30.133.1 0.0.0.255 0.0.0.0 255.255.255.255 30 deny ip 10.30.155.1 0.0.0.255 0.0.0.0 255.255.255.255 40 remark "THIS IS THE FINAL ACE IN THE LIST" 40 permit ip 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255 exit</pre>	<p>←</p> <p>←</p> <p>Removes an existing ACL and replaces it with a new version with the same identity. To append new ACEs to an existing ACL instead of replacing it, you would omit the first line and ensure that the sequence numbering for the new ACEs begin with a number greater than the highest number in the existing list.</p>
--	--

Figure 10-38. Example of an Offline ACL File Designed To Replace An Existing ACL

3. Use **copy tftp command-file** to download the file as a list of commands to the switch.

Example of Using the Offline Process

For example, suppose that you wanted to create an extended ACL for an RACL application to fulfill the following requirements (Assume a subnet mask of 255.255.255.0 and a TFTP server at 10.10.10.1.):

- ID: "LIST-20-IN"
- Deny Telnet access to a server at 10.10.10.100 on VLAN 10 from these three IP addresses on VLAN 20 (with ACL logging):
 - 10.10.20.17
 - 10.10.20.23
 - 10.10.20.40
- Allow any access to the server from all other addresses on VLAN 20:
- Permit internet access to these two IP address on VLAN 20, but deny access to all other addresses on VLAN 20 (without ACL logging).
 - 10.10.20.98
 - 10.10.20.21

- Deny all other IP traffic from VLAN 20 to VLAN 10.
 - Deny all IP traffic from VLAN 30 (10.10.30.0) to the server at 10.10.10.100 on VLAN 10 (without ACL logging), but allow any other IP traffic from VLAN 30 to VLAN 10.
 - Deny all other inbound IP traffic to VLAN 20. (Hint: The Implicit Deny can achieve this objective.)
1. You would create a **.txt** file with the content shown in figure 10-40.

```
ip access-list extended LIST-20-IN

; CREATED ON JUNE 27

10 remark "THIS ACE APPLIES INBOUND ON VLAN 20"
10 permit tcp any host 10.10.20.98 eq http
20 permit tcp any host 10.10.20.21 eq http
30 deny tcp any 10.10.20.1/24 eq http

; VLAN 20 SOURCES TO VLAN 10 DESTINATIONS.

40 deny tcp host 10.10.20.17 host 10.10.10.100 eq telnet log
50 deny tcp host 10.10.20.23 host 10.10.10.100 eq telnet log
60 deny tcp host 10.10.20.40 host 10.10.10.100 eq telnet log
70 permit ip 10.10.20.1/24 host 10.10.10.100
80 remark "VLAN 30 POLICY."
80 deny ip 10.10.30.1/24 host 10.10.10.100
90 permit ip 10.10.30.1/24 10.10.10.1/24
exit
vlan 20 ip access-group "LIST-20-in" in
```

The ";" enables a comment in the file.

Note: You can use the ";" character to denote a comment. The file stored on your TFTP server retains comments, and they appear when you use **copy** to download the ACL command file. (Comments are not saved in the switch configuration.)

Figure 10-39. Example of a .txt File Designed for Creating an ACL

2. After you copy the above .txt file to a TFTP server the switch can access, you would then execute the following command:

copy tftp command-file 10.10.10.1 LIST-20-IN.txt pc

In this example, the CLI would show the following output to indicate that the ACL was successfully downloaded to the switch:

Note

If a transport error occurs, the switch does not execute the command and the ACL is not configured.

```
ProCurve(config)# copy tftp command-file 10.10.10.1 LIST-20-IN.txt pc
Running configuration may change, do you want to continue [y/n]? Y
 1. ip access-list extended LIST-20-IN
 3. ; CREATED ON JUNE 27
 5. 10 remark "THIS ACE APPLIES INBOUND ON VLAN 20"
 6. 10 permit tcp any host 10.10.20.98 eq http
 7. 20 permit tcp any host 10.10.20.21 eq http
 8. 30 deny tcp any 10.10.20.1/24 eq http
10. ; VLAN 20 SOURCES TO VLAN 10 DESTINATIONS.
12. 40 deny tcp host 10.10.20.17 host 10.10.10.100 eq telnet log
13. 50 deny tcp host 10.10.20.23 host 10.10.10.100 eq telnet log
14. 60 deny tcp host 10.10.20.40 host 10.10.10.100 eq telnet log
15. 70 permit ip 10.10.20.1/24 host 10.10.10.100
16. 80 remark "VLAN 30 POLICY."
17. 80 deny ip 10.10.30.1/24 host 10.10.10.100
18. 90 permit ip 10.10.30.1/24 10.10.10.1/24
19. exit
20. vlan 20 ip access-group "LIST-20-in" in
```

As illustrated here, blank lines in the .txt file in figure 10-38 cause breaks in the displayed line-numbering sequence when you copy the command file to the switch. This is normal operation. (See also figure 10-41 for the configuration resulting from this output.)

Figure 10-40. Example of Using “copy tftp command-file” To Configure an ACL in the Switch

3. In this example, the command to assign the ACL to a VLAN was included in the .txt command file. If this is not done in your applications, then the next step is to manually assign the new ACL to the intended VLAN.

vlan < vid > ip access-group < identifier > in

4. You can then use the **show run** or **show access-list config** command to inspect the switch configuration to ensure that the ACL was properly downloaded.

Access Control Lists (ACLs) Creating or Editing ACLs Offline

```
ProCurve(config)# show run
. . .
ip access-list extended "LIST-20-IN"
 10 remark "THIS ACE APPLIES INBOUND ON VLAN 20"
 10 permit tcp 0.0.0.0 255.255.255.255 10.10.20.98 0.0.0.0 eq 80
 20 permit tcp 0.0.0.0 255.255.255.255 10.10.20.21 0.0.0.0 eq 80
 30 deny tcp 0.0.0.0 255.255.255.255 10.10.20.1 0.0.0.255 eq 80
 40 deny tcp 10.10.20.17 0.0.0.0 10.10.10.100 0.0.0.0 eq 23 log
 50 deny tcp 10.10.20.23 0.0.0.0 10.10.10.100 0.0.0.0 eq 23 log
 60 deny tcp 10.10.20.40 0.0.0.0 10.10.10.100 0.0.0.0 eq 23 log
 70 permit ip 10.10.20.1 0.0.0.255 10.10.10.100 0.0.0.0
 80 remark "VLAN 30 POLICY."
 80 deny ip 10.10.30.1 0.0.0.255 10.10.10.100 0.0.0.0
 90 permit ip 10.10.30.1 0.0.0.255 10.10.10.1 0.0.0.255
  exit
. . .
vlan 20
  name "VLAN20"
  no ip address
  ip access-group "LIST-20-in" in
  exit
```

Note that the comments preceded by ";" in the .txt source file for this configuration do not appear in the ACL configured in the switch.

As a part of the instruction set included in the .txt file, the ACL is assigned to inbound IP traffic on VLAN 20.

Figure 10-41. Example of Verifying the .txt File Download to the Switch

5. 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 IP traffic you do not want forwarded
- Receive notification when the switch detects attempts to forward IP traffic you have designed your ACLs to reject (deny)

The switch sends ACL messages to Syslog and optionally to the current console, Telnet, or SSH session. You can use **logging < >** to configure up to six Syslog server destinations.

Requirements for Using ACL Logging

- The switch configuration must include an ACL (1) assigned to a port, trunk, or static VLAN interface and (2) containing an ACE configured with the **deny** action and the **log** option.
- If the RACL application is used, then IP routing must be enabled on the switch.
- For ACL logging to a Syslog server:
 - The server must be accessible to the switch and identified in the running configuration.
 - The logging facility must be enabled for Syslog.
 - Debug must be configured to:
 - support ACL messages
 - send debug messages to the desired debug destination

These requirements are described in more detail under “Enabling ACL Logging on the Switch” on page 10-111.

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 10-42.




Figure 10-42. Content of a Message Generated by an ACL-Deny Action

Enabling ACL Logging on the Switch

1. If you are using a Syslog server, use the **logging < ip-addr >** command to configure the Syslog server IP address(es). Ensure that the switch can access any Syslog server(s) you specify.
2. Use **logging facility syslog** to enable the logging for Syslog operation.
3. Use the **debug destination** command to configure one or more log destinations. (Destination options include **logging**, **session**, and **windshell**. For more information on debug, refer to "Debug and Syslog Messaging Operation" in appendix C, "Troubleshooting", in the *Management and Configuration Guide* for your switch.)
4. Use **debug acl** or **debug all** to configure the debug operation to include ACL messages.
5. Configure one or more ACLs with the **deny** action and the **log** option.

For example, suppose that you want to configure the following operation:

- On VLAN 10 configure an extended ACL with an ACL-ID of "NO-TELNET" and use the **RACL in** option to deny Telnet traffic entering the switch from IP address 10.10.10.3 to any routed destination. (Note that this assignment will not filter Telnet traffic from 10.10.10.3 to destinations on VLAN 10 itself.)
- Configure the switch to send an ACL log message to the current console session and to a Syslog server at IP address 10.10.20.3 on VLAN 20 if the switch detects a packet match denying a Telnet attempt from 10.10.10.3.

(This example assumes that IP routing is already configured on the switch.)




Figure 10-43. Example of an ACL Log Application

Access Control Lists (ACLs) Enable ACL "Deny" Logging

```
ProCurve(config)# ip access-list extended NO-TELNET
ProCurve(config-ext-nacl)# remark "DENY 10.10.10.3 TELNET TRAFFIC IN"
ProCurve(config-ext-nacl)# deny tcp host 10.10.10.3 any eq telnet log
ProCurve(config-ext-nacl)# permit ip any any
ProCurve(config-ext-nacl)# exit
ProCurve(config)# vlan 10 ip access-group NO-TELNET in
ProCurve(config)# logging 10.10.20.3
ProCurve(config)# logging facility syslog
ProCurve(config)# debug destination logging
ProCurve(config)# debug destination session
ProCurve(config)# debug acl
ProCurve(config)# write mem
ProCurve(config)# show debug

Debug Logging

Destination:
Logging --
 10.10.20.3
  Facility = syslog
Session

Enabled debug types:
event
acl log

ProCurve(config)# show access-list config

ip access-list extended "NO-TELNET"
 10 remark "DENY 10.10.10.3 TELNET TRAFFIC"
 10 deny tcp 10.10.10.5 0.0.0.0 0.0.0.0 255.255.255.255 eq 23 log
 20 permit ip 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255
exit
```



Figure 10-44. Commands for Applying an ACL with Logging to Figure 10-43

General ACL Operating Notes

ACLs do not provide DNS hostname support. ACLs cannot be configured to screen hostname IP traffic between the switch and a DNS.

ACLs Do Not Affect Serial Port Access. ACLs do not apply to the switch's serial port.

ACL Screening of IP Traffic Generated by the Switch. Outbound RACL applications on a switch do not screen IP traffic (such as broadcasts, Telnet, Ping, and ICMP replies) *generated by the switch itself*. Note that all ACLs applied on the switch do screen this type of IP traffic when other devices generate it. Similarly, all ACL applications can screen responses from other devices to unscreened IP traffic the switch generates.

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 the last entry in an ACL as an explicit **deny** statement with a **log** statement included, and apply the ACL to an appropriate VLAN.
- 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 resource use by configuring the appropriate ACEs to match with specific hosts instead of entire subnets. (For more on resource usage, refer to “Monitoring Shared Resources” on page 10-114.)

Minimum Number of ACEs in an ACL. Any ACL must include at least one ACE to enable IP traffic screening. A numbered ACL cannot be created without at least one ACE. A named ACL can be created “empty”; that is, without any ACEs. However in an empty ACL applied to an interface, the Implicit Deny function does not operate, and the ACL has no effect on traffic.

Monitoring Shared Resources. Applied ACLs share internal switch resources with several other features. The switch provides ample resources for all features. However, if the internal resources become fully subscribed, additional ACLs cannot be applied until the necessary resources are released from other applications. For information on determining current resource availability and usage, refer to appendix E, “Monitoring Resources” in the *Management and Configuration Guide* for your switch.

Protocol Support . ACL criteria does not include use of MAC information or QoS.

Replacing or Adding To an Active ACL Policy. If you assign an ACL to an interface and subsequently add or replace ACEs in that ACL, each new ACE becomes active when you enter it. If the ACL is configured on multiple interfaces when the change occurs, then the switch resources must accommodate all applications of the ACL. If there are insufficient resources to accommodate one of several ACL applications affected by the change, then the change is not applied to any of the interfaces and the previous version of the ACL remains in effect. Refer to “Monitoring Shared Resources”, above.

“Strict” TCP and UDP. When the ACL configuration includes TCP or UDP options, the switch operates in “strict” TCP and UDP mode for increased control. In this case, the switch compares all TCP and UDP packets against the ACLs. (In the ProCurve 9300m and 9404sl Routing Switches, the Strict TCP and Strict UDP modes are optional and must be specifically invoked.)

Configuring Advanced Threat Protection

Contents

Introduction	11-2
DHCP Snooping	11-3
Overview	11-3
Enabling DHCP Snooping	11-4
Enabling DHCP Snooping on VLANs	11-6
Configuring DHCP Snooping Trusted Ports	11-7
Configuring Authorized Server Addresses	11-8
Using DHCP Snooping with Option 82	11-8
Changing the Remote-id from a MAC to an IP Address	11-10
Disabling the MAC Address Check	11-10
The DHCP Binding Database	11-11
Operational Notes	11-12
Log Messages	11-13
Dynamic ARP Protection	11-15
Introduction	11-15
Enabling Dynamic ARP Protection	11-17
Configuring Trusted Ports	11-17
Adding an IP-to-MAC Binding to the DHCP Database	11-18
Configuring Additional Validation Checks on ARP Packets	11-19
Verifying the Configuration of Dynamic ARP Protection	11-20
Displaying ARP Packet Statistics	11-21
Monitoring Dynamic ARP Protection	11-21
Using the Instrumentation Monitor	11-22
Operating Notes	11-23
Configuring Instrumentation Monitor	11-24
Examples	11-25
Viewing the Current Instrumentation Monitor Configuration	11-26

Introduction

As your network expands to include an increasing number of mobile devices, continuous Internet access, and new classes of users (such as partners, temporary employees, and visitors), additional protection from attacks launched from both inside and outside your internal network is often necessary.

Advanced threat protection can detect port scans and hackers who try to access a port or the switch itself. The following software features provide advanced threat protection and are described in this chapter:

- DHCP snooping: Protects your network from common DHCP attacks, such as:
 - Address spoofing in which an invalid IP address or network gateway address is assigned by a rogue DHCP server.
 - Address exhaustion of available addresses in the network DHCP server, caused by repeated attacker access to the network and numerous IP address requests.
- Dynamic ARP protection: Protects your network from ARP cache poisoning as in the following cases:
 - An unauthorized device forges an illegitimate ARP response and network devices use the response to update their ARP caches.
 - A denial-of-service (DoS) attack from unsolicited ARP responses changes the network gateway IP address so that outgoing traffic is prevented from leaving the network and overwhelms network devices.
- Instrumentation monitor: Protects your network from a variety of other common attacks besides DHCP and ARP attacks, including:
 - Attempts at a port scan to expose a vulnerability in the switch, indicated by an excessive number of packets sent to closed TCP/UDP ports
 - Attempts to fill all IP address entries in the switch's forwarding table and cause legitimate traffic to be dropped, indicated by an increased number of learned IP destination addresses
 - Attempts to spread viruses, indicated by an increased number of ARP request packets

- Attempts to exhaust system resources so that sufficient resources are not available to transmit legitimate traffic, indicated by an unusually high use of specific system resources
- Attempts to attack the switch's CPU and introduce delay in system response time to new network events
- Attempts by hackers to access the switch, indicated by an excessive number of failed logins or port authentication failures
- Attempts to deny switch service by filling the forwarding table, indicated by an increased number of learned MAC addresses or a high number of MAC address moves from one port to another
- Attempts to exhaust available CPU resources, indicated by an increased number of learned MAC address events being discarded

DHCP Snooping

Command	Page
dhcp-snooping	page 11-4
authorized-server	page 11-8
database	page 11-11
option	page 11-8
trust	page 11-7
verify	page 11-10
vlan	page 11-6
show dhcp-snooping	page 11-5
show dhcp-snooping stats	page 11-5
dhcp-snooping binding	page 11-12
debug dhcp-snooping	page 11-12

Overview

You can use DHCP snooping to help avoid the Denial of Service attacks that result from unauthorized users adding a DHCP server to the network that then provides invalid configuration data to other DHCP clients on the network.

DHCP snooping accomplishes this by allowing you to distinguish between trusted ports connected to a DHCP server or switch and untrusted ports connected to end-users. DHCP packets are forwarded between trusted ports without inspection. DHCP packets received on other switch ports are inspected before being forwarded. Packets from untrusted sources are dropped. Conditions for dropping packets are shown below.

Condition for Dropping a Packet	Packet Types
A packet from a DHCP server received on an untrusted port	DHCPOFFER, DHCPACK, DHCPNACK
If the switch is configured with a list of authorized DHCP server addresses and a packet is received from a DHCP server on a trusted port with a source IP address that is not in the list of authorized DHCP server addresses.	DHCPOFFER, DHCPACK, DHCPNACK
Unless configured to not perform this check, a DHCP packet received on an untrusted port where the DHCP client hardware address field does not match the source MAC address in the packet	N/A
Unless configured to not perform this check, a DHCP packet containing DHCP relay information (option 82) received from an untrusted port	N/A
A broadcast packet that has a MAC address in the DHCP binding database, but the port in the DHCP binding database is different from the port on which the packet is received	DHCPRELEASE, DHCPDECLINE

Enabling DHCP Snooping

DHCP snooping is enabled globally by entering this command:

```
ProCurve (config) # dhcp-snooping
```

Use the **no** form of the command to disable DHCP snooping.

Syntax: [no] dhcp-snooping [authorized-server | database | option | trust | verify | vlan]

authorized server: Enter the IP address of a trusted DHCP server. If no authorized servers are configured, all DHCP server addresses are considered valid.
Maximum: 20 authorized servers

database: To configure a location for the lease database, enter a URL in the format **tftp://ip-addr/ascii-string**. The maximum number of characters for the URL is 63.

option: Add relay information option (Option 82) to DHCP client packets that are being forwarded out trusted ports. The default is **yes**, add relay information.

trust: Configure trusted ports. Only server packets received on trusted ports are forwarded. Default: **untrusted**.

verify: Enables DHCP packet validation. The DHCP client hardware address field and the source MAC address must be the same for packets received on untrusted ports or the packet is dropped. Default: **Yes**

vlan: Enable DHCP snooping on a vlan. DHCP snooping must be enabled already. Default: **No**

To display the DHCP snooping configuration, enter this command:

```
ProCurve(config)# show dhcp-snooping
```

An example of the output is shown below.

```
ProCurve(config)# show dhcp-snooping
DHCP Snooping Information
  DHCP Snooping           : Yes
  Enabled Vlans           :
  Verify MAC              : Yes
  Option 82 untrusted policy : drop
  Option 82 Insertion     : Yes
  Option 82 remote-id     : mac
  Store lease database    : Not configured
  Port Trust
  -----
  B1      No
  B2      No
  B3      No
```

Figure 11-1. An Example of the DHCP Snooping Command Output

To display statistics about the DHCP snooping process, enter this command:

```
ProCurve(config)# show dhcp-snooping stats
```

An example of the output is shown below.

```
ProCurve(config)# show dhcp-snooping stats
```

Packet type	Action	Reason	Count
server	forward	from trusted port	8
client	forward	to trusted port	8
server	drop	received on untrusted port	2
server	drop	unauthorized server	0
client	drop	destination on untrusted port	0
client	drop	untrusted option 82 field	0
client	drop	bad DHCP release request	0
client	drop	failed verify MAC check	0

Figure 11-2. Example of Show DHCP Snooping Statistics

Enabling DHCP Snooping on VLANs

DHCP snooping on VLANs is disabled by default. To enable DHCP snooping on a VLAN or range of VLANs enter this command:

```
ProCurve(config)# dhcp-snooping vlan <vlan-id-range>
```

You can also use this command in the vlan context, in which case you cannot enter a range of VLANs for snooping.

Below is an example of DHCP snooping enabled on VLAN 4.

```
ProCurve(config)# dhcp-snooping vlan 4
ProCurve(config)# show dhcp-snooping
```

DHCP Snooping Information

DHCP Snooping	: Yes
Enabled Vlans	: 4
Verify MAC	: Yes
Option 82 untrusted policy	: drop
Option 82 Insertion	: Yes
Option 82 remote-id	: mac

Figure 11-3. Example of DHCP Snooping on a VLAN

Configuring DHCP Snooping Trusted Ports

By default, all ports are untrusted. To configure a port or range of ports as trusted, enter this command:

```
ProCurve(config)# dhcp-snooping trust <port-list>
```

You can also use this command in the interface context, in which case you are not able to enter a list of ports.

```
ProCurve(config)# dhcp-snooping trust B1-B2
ProCurve(config)# show dhcp-snooping

DHCP Snooping Information

DHCP Snooping           : Yes
Enabled Vlans           : 4
Verify MAC              : Yes
Option 82 untrusted policy : drop
Option 82 Insertion     : Yes
Option 82 remote-id     : mac

Store lease database : Not configured

Port  Trust
-----
B1    Yes
B2    Yes
B3    No
```

Figure 11-4. Example of Setting Trusted Ports

DHCP server packets are forwarded only if received on a trusted port; DHCP server packets received on an untrusted port are dropped.

Use the **no** form of the command to remove the trusted configuration from a port.

Configuring Authorized Server Addresses

If authorized server addresses are configured, a packet from a DHCP server must be received on a trusted port AND have a source address in the authorized server list in order to be considered valid. If no authorized servers are configured, all servers are considered valid. You can configure a maximum of 20 authorized servers.

To configure a DHCP authorized server address, enter this command in the global configuration context:

```
ProCurve(config)# dhcp-snooping authorized-server  
                <ip-address>
```

```
ProCurve(config)# show dhcp-snooping

DHCP Snooping Information

  DHCP Snooping           : Yes
  Enabled Vlans           : 4
  Verify MAC              : No
  Option 82 untrusted policy : drop
  Option 82 Insertion     : Yes
  Option 82 remote-id     : subnet-ip

Authorized Servers
-----
111.222.3.4
```

Figure 11-5. Example of Authorized Servers for DHCP Snooping

Using DHCP Snooping with Option 82

DHCP adds Option 82 (relay information option) to DHCP request packets received on untrusted ports by default. (See the preceding section *Configuring DHCP Relay* for more information on Option 82.)

When DHCP is enabled globally and also enabled on a VLAN, and the switch is acting as a DHCP relay, the settings for the DHCP relay Option 82 command are ignored when snooping is controlling Option 82 insertion. Option 82 inserted in this manner allows the association of the client's lease with the correct port, even when another device is acting as a DHCP relay or when the server is on the same subnet as the client.

Note

DHCP snooping only overrides the Option 82 settings on a VLAN that has snooping enabled, not on VLANS without snooping enabled.

If DHCP snooping is enabled on a switch where an edge switch is also using DHCP snooping, it is desirable to have the packets forwarded so the DHCP bindings are learned. To configure the policy for DHCP packets from untrusted ports that already have Option 82 present, enter this command in the global configuration context.

Syntax: [no] dhcp-snooping option 82 [remote-id <mac | subnet-ip | mgmt-ip>]
[untrusted-policy <drop | keep | replace>]

Enables DHCP Option 82 insertion in the packet.

remote-id *Set the value used for the **remote-id** field of the relay information option.*

mac: *The switch mac address is used for the remote-id. This is the default.*

subnet-ip: *The IP address of the VLAN the packet was received on is used for the remote-id. If **subnet-ip** is specified but the value is not set, the MAC address is used.*

mgmt-ip: *The management VLAN IP address is used as the remote-id. If **mgmt-ip** is specified but the value is not set, the MAC address is used.*

untrusted-policy *Configures DHCP snooping behavior when forwarding a DHCP packet from an untrusted port that already contains DHCP relay information (Option 82). The default is **drop**.*

drop: *The packet is dropped.*

keep: *The packet is forwarded without replacing the option information.*

replace: *The existing option is replaced with a new Option 82 generated by the switch.*

Note

The default **drop** policy should remain in effect if there are any untrusted nodes, such as clients, directly connected to this switch.

Changing the Remote-id from a MAC to an IP Address

By default, DHCP snooping uses the MAC address of the switch as the remote-id in Option 82 additions. The IP address of the VLAN the packet was received on or the IP address of the management VLAN can be used instead by entering this command with the associated parameter:

```
ProCurve(config)# dhcp-snooping option 82 remote-id  
                  <mac|subnet-ip|mgmt-ip>
```

```
ProCurve(config)# dhcp-snooping option 82 remote-id subnet-  
ip  
ProCurve(config)# show dhcp-snooping  
  
DHCP Snooping Information  
  
DHCP Snooping           : Yes  
Enabled Vlans           : 4  
Verify MAC               : Yes  
Option 82 untrusted policy : drop  
Option 82 Insertion      : Yes  
Option 82 remote-id      : subnet-ip
```

Figure 11-6. Example of DHCP Snooping Option 82 using the VLAN IP Address

Disabling the MAC Address Check

DHCP snooping drops DHCP packets received on untrusted ports when the check address (chaddr) field in the DHCP header does not match the source MAC address of the packet (default behavior). To disable this checking, use the **no** form of this command.

```
ProCurve(config)# dhcp-snooping verify mac
```

```
ProCurve(config)# dhcp-snooping verify mac
ProCurve(config)# show dhcp-snooping

DHCP Snooping Information

DHCP Snooping           : Yes
Enabled Vlans           : 4
Verify MAC               : yes
Option 82 untrusted policy : drop
Option 82 Insertion     : Yes
Option 82 remote-id     : subnet-ip
```

Figure 11-7. Example Showing the DHCP Snooping Verify MAC Setting

The DHCP Binding Database

DHCP snooping maintains a database of up to 8192 DHCP bindings on untrusted ports. Each binding consists of:

- Client MAC address
- Port number
- VLAN identifier
- Leased IP address
- Lease time

The switch can be configured to store the bindings at a specific URL so they will not be lost if the switch is rebooted. If the switch is rebooted, it will read its binding database from the specified location. To configure this location use this command.

Syntax: [no] dhcp-snooping database [file<tftp://<ip-address>/<ascii-string>>]
[delay<15-86400>][timeout<0-86400>]

- | | |
|----------------|---|
| file | <i>Must be in Uniform Resource Locator (URL) format — “tftp://ip-address/ascii-string”. The maximum filename length is 63 characters.</i> |
| delay | <i>Number of seconds to wait before writing to the database. Default = 300 seconds.</i> |
| timeout | <i>Number of seconds to wait for the database file transfer to finish before returning an error. A value of zero (0) means retry indefinitely. Default = 300 seconds.</i> |

A message is logged in the system event log if the DHCP binding database fails to update.

To display the contents of the DHCP snooping binding database, enter this command.

Syntax: show dhcp-snooping binding

```
ProCurve(config)# show dhcp-snooping binding
```

MacAddress	IP	VLAN	Interface	Time left
22.22.22.22.22.22	10.0.0.1	4	B2	1600

Figure 11-8. Example Showing DHCP Snooping Binding Database Contents

Note

If a lease database is configured, the switch drops all DHCP packets until the lease database is read. This only occurs when the switch reboots and is completed quickly. If the switch is unable to read the lease database from the tftp server, it waits until that operation times out and then begins forwarding DHCP packets.

Enabling Debug Logging

To enable debug logging for DHCP snooping, use this command.

Syntax: [no] debug dhcp-snooping [agent | event | packet]

agent *Displays DHCP snooping agent messages.*

event *Displays DHCP snooping event messages.*

packet *Displays DHCP snooping packet messages.*

Operational Notes

- DHCP is not configurable from the web management interface or menu interface.
- If packets are received at too high a rate, some may be dropped and need to be re-transmitted.

- ProCurve recommends running a time synchronization protocol such as SNTP in order to track lease times accurately.
- A remote server must be used to save lease information or there may be a loss of connectivity after a switch reboot.

Log Messages

Server <ip-address> packet received on untrusted port <port-number> dropped. Indicates a DHCP server on an untrusted port is attempting to transmit a packet. This event is recognized by the reception of a DHCP server packet on a port that is configured as untrusted.

Ceasing untrusted server logs for %s. More than one packet was received from a DHCP server on an untrusted port. To avoid filling the log file with repeated attempts, untrusted server drop packet events will not be logged for the specified <duration>.

Client packet destined to untrusted port <port-number> dropped. Indicates that the destination of a DHCP client unicast packet is on an untrusted port. This event is recognized when a client unicast packet is dropped because the destination address is out a port configured as untrusted.

Ceasing untrusted port destination logs for %s. More than one client unicast packet with an untrusted port destination was dropped. To avoid filling the log file with repeated attempts, untrusted port destination attempts will not be logged for the specified <duration>.

Unauthorized server <ip-address> detected on port <port-number>. Indicates that an unauthorized DHCP server is attempting to send packets. This event is recognized when a server packet is dropped because there are configured authorized servers and a server packet is received from a server that is not configured as an authorized server.

Ceasing unauthorized server logs for <duration>. More than one unauthorized server packet was dropped. To avoid filling the log file with repeated attempts, unauthorized server transmit attempts will not be logged for the specified <duration>.

Received untrusted relay information from client <mac-address> on port <port-number>. Indicates the reception on an untrusted port of a client packet containing a relay information option field. This event is recognized when a client packet containing a relay information option field is dropped because it was received on a port configured as untrusted.

Ceasing untrusted relay information logs for <duration>. More than one DHCP client packet received on an untrusted port with a relay information field was dropped. To avoid filling the log file with repeated attempts, untrusted relay information packets will not be logged for the specified <duration>.

Client address <mac-address> not equal to source MAC <mac-address> detected on port <port-number>. Indicates that a client packet source MAC address does not match the “chaddr” field. This event is recognized when the dhcp-snooping agent is enabled to filter DHCP client packets that do not have a matching “chaddr” field and source MAC address.

Ceasing MAC mismatch logs for <duration>. More than one DHCP client packet with a mismatched source MAC and chaddr field was dropped. To avoid filling the log file with repeated attempts, client address mismatch events will not be logged for the specified <duration>.

Attempt to release address <ip-address> leased to port <port-number> detected on port <port-number> dropped. Indicates an attempt by a client to release an address when a DHCPRELEASE or DHCPDECLINE packet is received on a port different from the port the address was leased to.

Ceasing bad release logs for %s. More than one bad DHCP client release packet was dropped. To avoid filling the log file with repeated bad release dropped packets, bad releases will not be logged for <duration>.

Lease table is full, DHCP lease was not added. The lease table is full and this lease will not be added to it.

Write database to remote file failed errno (error-num). An error occurred while writing the temporary file and sending it using tftp to the remote server.

DHCP packets being rate-limited. Too many DHCP packets are flowing through the switch and some are being dropped.

Snooping table is full. The DHCP binding table is full and subsequent bindings are being dropped.

Dynamic ARP Protection

Introduction

On the VLAN interfaces of a routing switch, dynamic ARP protection ensures that only valid ARP requests and responses are relayed or used to update the local ARP cache. ARP packets with invalid IP-to-MAC address bindings advertised in the source protocol address and source physical address fields are discarded. For more information about the ARP cache, refer to “ARP Cache Table” in the *Multicast and Routing Guide*.

ARP requests are ordinarily broadcast, and received by all devices in a broadcast domain. Most ARP devices update their IP-to-MAC address entries each time they receive an ARP packet even if they did not request the information. This behavior makes an ARP cache vulnerable to attacks.

Because ARP allows a node to update its cache entries on other systems by broadcasting or unicasting a gratuitous ARP reply, an attacker can send his own IP-to-MAC address binding in the reply that causes all traffic destined for a VLAN node to be sent to the attacker's MAC address. As a result, the attacker can intercept traffic for other hosts in a classic "man-in-the-middle" attack. The attacker gains access to any traffic sent to the poisoned address and can capture passwords, e-mail, and VoIP calls or even modify traffic before resending it.

Another way in which the ARP cache of known IP addresses and associated MAC addresses can be poisoned is through unsolicited ARP responses. For example, an attacker can associate the IP address of the network gateway with the MAC address of a network node. In this way, all outgoing traffic is prevented from leaving the network because the node does not have access to outside networks. As a result, the node is overwhelmed by outgoing traffic destined to another network.

Dynamic ARP protection is designed to protect your network against ARP poisoning attacks in the following ways:

- Allows you to differentiate between trusted and untrusted ports.
- Intercepts all ARP requests and responses on untrusted ports before forwarding them.
- Verifies IP-to-MAC address bindings on untrusted ports with the information stored in the lease database maintained by DHCP snooping and user-configured static bindings (in non-DHCP environments):

- If a binding is valid, the switch updates its local ARP cache and forwards the packet.
- If a binding is invalid, the switch drops the packet, preventing other network devices from receiving the invalid IP-to-MAC information.

DHCP snooping intercepts and examines DHCP packets received on switch ports before forwarding the packets. DHCP packets are checked against a database of DHCP binding information. Each binding consists of a client MAC address, port number, VLAN identifier, leased IP address, and lease time. The DHCP binding database is used to validate packets by other security features on the switch. For more information, refer to “DHCP Snooping” in the *Access Security Guide*.

If you have already enabled DHCP snooping on a switch, you may also want to add static IP-to-MAC address bindings to the DHCP snooping database so that ARP packets from devices that have been assigned static IP addresses are also verified.

- Supports additional checks to verify source MAC address, destination MAC address, and IP address.

ARP packets that contain invalid IP addresses or MAC addresses in their body that do not match the addresses in the Ethernet header are dropped.

When dynamic ARP protection is enabled, only ARP request and reply packets with valid IP-to-MAC address bindings in their packet header are relayed and used to update the ARP cache.

Dynamic ARP protection is implemented in the following ways on a switch:

- You can configure dynamic ARP protection only from the CLI; you cannot configure this feature from the web or menu interfaces.
- Line rate—Dynamic ARP protection copies ARP packets to the switch CPU, evaluates the packets, and then re-forwards them through the switch software. During this process, if ARP packets are received at too high a line rate, some ARP packets may be dropped and will need to be retransmitted.
- The SNMP MIB, HP-ICF-ARP-PROTECT-MIB, is created to configure dynamic ARP protection and to report ARP packet-forwarding status and counters.

Enabling Dynamic ARP Protection

To enable dynamic ARP protection for VLAN traffic on a routing switch, enter the **arp protect vlan** command at the global configuration level.

Syntax: [no] arp protect vlan {vlan-range}

vlan-range Specifies a VLAN ID or a range of VLAN IDs from one to 4094; for example, 1–200.

An example of the **arp protect vlan** command is shown here:

```
ProCurve(config)# arp protect vlan 1-101
```

Configuring Trusted Ports

In a similar way to DHCP snooping, dynamic ARP protection allows you to configure VLAN interfaces in two categories: trusted and untrusted ports. ARP packets received on trusted ports are forwarded without validation.

By default, all ports on a switch are untrusted. If a VLAN interface is untrusted:

- The switch intercepts all ARP requests and responses on the port.
- Each intercepted packet is checked to see if its IP-to-MAC binding is valid. If a binding is invalid, the switch drops the packet.

You must configure trusted ports carefully. For example, in the topology in Figure 11-9, Switch B may not see the leased IP address that Host 1 receives from the DHCP server. If the port on Switch B that is connected to Switch A is untrusted and if Switch B has dynamic ARP protection enabled, it will see ARP packets from Host 1 as invalid, resulting in a loss of connectivity.

On the other hand, if Switch A does not support dynamic ARP protection and you configure the port on Switch B connected to Switch A as trusted, Switch B opens itself to possible ARP poisoning from hosts attached to Switch A.




Figure 11-9. Configuring Trusted Ports for Dynamic ARP Protection

Take into account the following configuration guidelines when you use dynamic ARP protection in your network:

- You should configure ports connected to other switches in the network as trusted ports. In this way, all network switches can exchange ARP packets and update their ARP caches with valid information.
- Switches that do not support dynamic ARP protection should be separated by a router in their own Layer 2 domain. Because ARP packets do not cross Layer 2 domains, the unprotected switches cannot unknowingly accept ARP packets from an attacker and forward them to protected switches through trusted ports.

To configure one or more Ethernet interfaces that handle VLAN traffic as trusted ports, enter the **arp protect trust** command at the global configuration level. The switch does not check ARP requests and responses received on a trusted port.

Syntax: [no] arp protect trust <port-list>

<i>port-list</i>	Specifies a port number or a range of port numbers. Separate individual port numbers or ranges of port numbers with a comma; for example: c1-c3, c6 .
------------------	--

An example of the **arp protect trust** command is shown here:

```
ProCurve(config)# arp protect trust b1-b4, d1
```

Adding an IP-to-MAC Binding to the DHCP Database

A routing switch maintains a DHCP binding database, which is used for DHCP and ARP packet validation. Both the DHCP snooping and DHCP Option 82 insertion features maintain the lease database by learning the IP-to-MAC bindings on untrusted ports. Each binding consists of the client MAC address, port number, VLAN identifier, leased IP address, and lease time.

If your network does not use DHCP or if some network devices have fixed, user-configured IP addresses, you can enter static IP-to-MAC address bindings in the DHCP binding database. The switch uses manually configured static bindings for DHCP snooping and dynamic ARP protection.

To add the static configuration of an IP-to-MAC binding for a port to the database, enter the **ip source binding** command at the global configuration level.

Syntax: [no] ip source binding <mac-address> vlan <vlan-id> <ip-address>
interface <port-number>

mac-address Specifies a MAC address to bind with a VLAN and IP address on the specified port in the DHCP binding database.

vlan vlan-id Specifies a VLAN ID number to bind with the specified MAC and IP addresses on the specified port in the DHCP binding database.

ip-address Specifies an IP address to bind with a VLAN and MAC address on the specified port in the DHCP binding database.

interface port-number Specifies the port number on which the IP-to-MAC address and VLAN binding is configured in the DHCP binding database.

An example of the **ip source binding** command is shown here:

```
ProCurve(config)# ip source binding 0030c1-7f49c0  
interface vlan 100 10.10.20.1 interface A4
```

Note

Note that the **ip source binding** command is the same command used by the Dynamic IP Lockdown feature to configure static bindings. The Dynamic ARP Protection and Dynamic IP Lockdown features share a common list of source IP-to-MAC bindings.

Configuring Additional Validation Checks on ARP Packets

Dynamic ARP protection can be configured to perform additional validation checks on ARP packets. By default, no additional checks are performed. To configure additional validation checks, enter the **arp protect validate** command at the global configuration level.

Syntax: [no] arp protect validate <[src-mac] | [dst-mac] | [ip]>

src-mac (Optional) Drops any ARP request or response packet in which the source MAC address in the Ethernet header does not match the sender MAC address in the body of the ARP packet.

- dst-mac** *(Optional) Drops any unicast ARP response packet in which the destination MAC address in the Ethernet header does not match the target MAC address in the body of the ARP packet.*
- ip** *(Optional) Drops any ARP packet in which the sender IP address is invalid. Drops any ARP response packet in which the target IP address is invalid. Invalid IP addresses include: 0.0.0.0, 255.255.255.255, all IP multicast addresses, and all Class E IP addresses.*

You can configure one or more of the validation checks. The following example of the **arp protect validate** command shows how to configure the validation checks for source MAC address and destination MAC address:

```
ProCurve(config)# arp protect validate src-mac dst-mac
```

Verifying the Configuration of Dynamic ARP Protection

To display the current configuration of dynamic ARP protection, including the additional validation checks and the trusted ports that are configured, enter the **show arp protect** command:

```
ProCurve(config)# show arp protect

ARP Protection Information

Enabled Vlans   : 1-4094
Validate       : dst-mac, src-mac

Port  Trust
-----
B1    Yes
B2    Yes
B3    No
B4    No
B5    No
```

Displaying ARP Packet Statistics

To display statistics about forwarded ARP packets, dropped ARP packets, MAC validation failure, and IP validation failures, enter the **show arp protect statistics** command:

```
ProCurve(config)# show arp protect statistics

Status and Counters - ARP Protection Counters for VLAN 1
Forwarded pkts      : 10      Bad source mac      : 2
Bad bindings        : 1       Bad destination mac : 1
Malformed pkts     : 0       Bad IP address      : 0

Status and Counters - ARP Protection Counters for VLAN 2
Forwarded pkts      : 1       Bad source mac      : 1
Bad bindings        : 1       Bad destination mac : 1
Malformed pkts     : 1       Bad IP address      : 1
```

Monitoring Dynamic ARP Protection

When dynamic ARP protection is enabled, you can monitor and troubleshoot the validation of ARP packets with the **debug arp protect** command. Use this command when you want to debug the following conditions:

- The switch is dropping valid ARP packets that should be allowed.
- The switch is allowing invalid ARP packets that should be dropped.

```
ProCurve(config)# debug arp protect

1. ARP request is valid
"DARPP: Allow ARP request 000000-000001,10.0.0.1 for 10.0.0.2 port A1,
vlan "

2. ARP request detected with an invalid binding
"DARPP: Deny ARP request 000000-000003,10.0.0.1 port A1, vlan 1"

3. ARP response with a valid binding
"DARPP: Allow ARP reply 000000-000002,10.0.0.2 port A2, vlan 1"

4. ARP response detected with an invalid binding
"DARPP: Deny ARP reply 000000-000003,10.0.0.2 port A2, vlan 1"
```

Using the Instrumentation Monitor

The instrumentation monitor can be used to detect anomalies caused by security attacks or other irregular operations on the switch. The following table shows the operating parameters that can be monitored at pre-determined intervals, and the possible security attacks that may trigger an alert:

Parameter Name	Description
pkts-to-closed-ports	The count of packets per minute sent to closed TCP/UDP ports. An excessive amount of packets could indicate a port scan, in which an attacker is attempting to expose a vulnerability in the switch.
arp-requests	The count of ARP requests processed per minute. A large amount of ARP request packets could indicate a host infected with a virus that is trying to spread itself.
ip-address-count	The number of destination IP addresses learned in the IP forwarding table. Some attacks fill the IP forwarding table causing legitimate traffic to be dropped.
system-resource-usage	The percentage of system resources in use. Some Denial-of-Service (DoS) attacks will cause excessive system resource usage, resulting in insufficient resources for legitimate traffic.

Parameter Name	Description
login-failures/min	The count of failed CLI login attempts or SNMP management authentication failures. This indicates an attempt has been made to manage the switch with an invalid login or password. Also, it might indicate a network management station has not been configured with the correct SNMP authentication parameters for the switch.
port-auth-failures/min	The count of times a client has been unsuccessful logging into the network
system-delay	The response time, in seconds, of the CPU to new network events such as BPDU packets or packets for other network protocols. Some DoS attacks can cause the CPU to take too long to respond to new network events, which can lead to a breakdown of Spanning Tree or other features. A delay of several seconds indicates a problem.
mac-address-count	The number of MAC addresses learned in the forwarding table. Some attacks fill the forwarding table so that new conversations are flooded to all parts of the network.
mac-moves/min	The average number of MAC address moves from one port to another per minute. This usually indicates a network loop, but can also be caused by DoS attacks.
learn-discards/min	Number of MAC address learn events per minute discarded to help free CPU resources when busy.

Operating Notes

- To generate alerts for monitored events, you must enable the instrumentation monitoring log and/or SNMP trap. The threshold for each monitored parameter can be adjusted to minimize false alarms (see “Configuring Instrumentation Monitor” on page 11-24).
- When a parameter exceeds its threshold, an alert (event log message and/or SNMP trap) is generated to inform network administrators of this condition. The following example shows an event log message that occurs when the number of MAC addresses learned in the forwarding table exceeds the configured threshold:

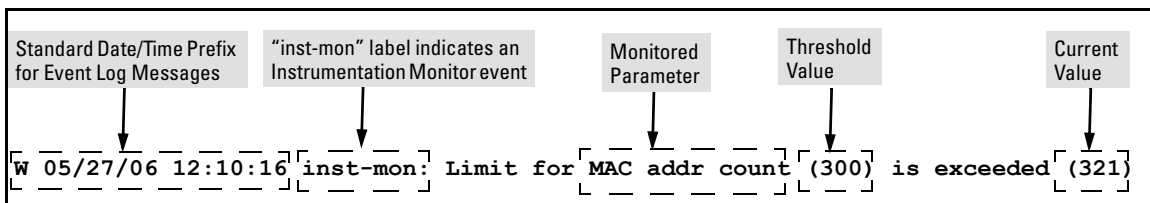


Figure 1. Example of Event Log Message generated by Instrumentation Monitor

- Alerts are automatically rate limited to prevent filling the log file with redundant information. The following is an example of alerts that occur when the device is continually subject to the same attack (too many MAC addresses in this instance):

```
W 01/01/90 00:05:00 inst-mon: Limit for MAC addr count (300) is exceeded (321)
W 01/01/90 00:10:00 inst-mon: Limit for MAC addr count (300) is exceeded (323)
W 01/01/90 00:15:00 inst-mon: Limit for MAC addr count (300) is exceeded (322)
W 01/01/90 00:20:00 inst-mon: Limit for MAC addr count (300) is exceeded (324)
W 01/01/90 00:20:00 inst-mon: Ceasing logs for MAC addr count for 15 minutes
```

Figure 2. Example of rate limiting when multiple messages are generated

In the preceding example, if a condition is reported 4 times (persists for more than 15 minutes) then alerts cease for 15 minutes. If after 15 minutes the condition still exists, the alerts cease for 30 minutes, then for 1 hour, 2 hours, 4 hours, 8 hours, and after that the persisting condition is reported once a day. As with other event log entries, these alerts can be sent to a syslog server.

- **Known Limitations:** The instrumentation monitor runs once every five minutes. The current implementation does not track information such as the port, MAC, and IP address from which an attack is received.

Configuring Instrumentation Monitor

The following commands and parameters are used to configure the operational thresholds that are monitored on the switch. By default, the instrumentation monitor is disabled.

Syntax: [no] instrumentation monitor [parameterName|all] [<low|med|high|limitValue>]

[log] : Enables/disables instrumentation monitoring log so that event log messages are generated every time there is an event which exceeds a configured threshold.

(Default threshold setting when instrumentation monitoring is enabled: **enabled**)

[all] : Enables/disables all counter types on the switch but does not enable/disable instrumentation monitor logging.

(Default threshold setting when enabled: **see parameter listings below**)

[arp-requests] : The number of arp requests that are processed each minute.

(Default threshold setting when enabled: **1000 (med)**)

[ip-address-count]: The number of destination IP addresses learned in the IP forwarding table.

(Default threshold setting when enabled: **1000 (med)**)

[learn-discards] : *The number of MAC address learn events per minute discarded to help free CPU resources when busy.*

*(Default threshold setting when enabled: **100 (med)**)*

[login-failures] : *The count of failed CLI login attempts or SNMP management authentication failures per hour.*

*(Default threshold setting when enabled: **10 (med)**)*

[mac-address-count] : *The number of MAC addresses learned in the forwarding table. You must enter a specific value in order to enable this feature.*

*(Default threshold setting when enabled: **1000 (med)**)*

[mac-moves] : *The average number of MAC address moves per minute from one port to another.*

*(Default threshold setting when enabled: **100 (med)**)*

[pkts-to-closed-ports] : *The count of packets per minute sent to closed TCP/UDP ports.*

*(Default threshold setting when enabled: **10 (med)**)*

[port-auth-failures] : *The count of times per minute that a client has been unsuccessful logging into the network.*

*(Default threshold setting when enabled: **10 (med)**)*

[system-resource-usage]: *The percentage of system resources in use.*

*(Default threshold setting when enabled: **50 (med)**)*

[system-delay] : *The response time, in seconds, of the CPU to new network events such as BPDUs packets or packets for other network protocols.*

*(Default threshold setting when enabled: **3 seconds (med)**)*

[trap] : *Enables or disables SNMP trap generation.*

*(Default setting when instrumentation monitoring is enabled: **disabled**)*

To enable instrumentation monitor using the default parameters and thresholds, enter the general **instrumentation monitor** command. To adjust specific settings, enter the name of the parameter that you wish to modify, and revise the threshold limits as needed.

Examples

To turn on monitoring and event log messaging with the default medium values:

```
ProCurve(config)# instrumentation monitor
```

To turn off monitoring of the system delay parameter:

```
ProCurve(config)# no instrumentation monitor system-  
delay
```

To adjust the alert threshold for the MAC address count to the low value:

```
ProCurve(config)# instrumentation monitor mac-  
address-count low
```

To adjust the alert threshold for the MAC address count to a specific value:

```
ProCurve(config)# instrumentation monitor mac-  
address-count 767
```

To enable monitoring of learn discards with the default medium threshold value:

```
ProCurve(config)# instrumentation monitor learn-  
discards
```

To disable monitoring of learn discards:

```
ProCurve(config)# no instrumentation monitor learn-  
discards
```

To enable or disable SNMP trap generation:

```
ProCurve(config)# [no] instrumentation monitor trap
```

Viewing the Current Instrumentation Monitor Configuration

The **show instrumentation monitor configuration** command displays the configured thresholds for monitored parameters.

```
ProCurve# show instrumentation monitor configuration
```

PARAMETER	LIMIT
mac-address-count	1000 (med)
ip-address-count	1000 (med)
system-resource-usage	50 (med)
system-delay	5 (high)
mac-moves/min	100 (med)
learn-discards/min	100 (med)
ip-port-scans/min	10 (med)
arp-requests/min	100 (low)
login-failures/min	10 (med)
port-auth-failures/min	10 (med)

```
SNMP trap generation for alerts: enabled  
Instrumentation monitoring log : enabled
```

Figure 11-3. Viewing the Instrumentation Monitor Configuration

An alternate method of determining the current Instrumentation Monitor configuration is to use the **show run** command. However, the show run command output does not display the threshold values for each limit set.

— This page is intentionally unused —

Traffic/Security Filters and Monitors

Contents

Overview	12-2
Introduction	12-2
Filter Limits	12-3
Using Port Trunks with Filters	12-3
Filter Types and Operation	12-3
Source-Port Filters	12-4
Operating Rules for Source-Port Filters	12-4
Example	12-5
Named Source-Port Filters	12-6
Operating Rules for Named Source-Port Filters	12-6
Defining and Configuring Named Source-Port Filters	12-7
Viewing a Named Source-Port Filter	12-8
Using Named Source-Port Filters	12-9
Static Multicast Filters	12-15
Protocol Filters	12-16
Configuring Traffic/Security Filters	12-17
Configuring a Source-Port Traffic Filter	12-18
Example of Creating a Source-Port Filter	12-19
Configuring a Filter on a Port Trunk	12-19
Editing a Source-Port Filter	12-20
Configuring a Multicast or Protocol Traffic Filter	12-21
Filter Indexing	12-22
Displaying Traffic/Security Filters	12-23

Overview

Applicable Switch Models. As of October, 2005, Traffic/Security filters are available on these current ProCurve switch models:

Switch Models	Source-Port Filters	Protocol Filters	Multicast Filters
Series 6400cl	Yes	No	No
Series 5400zl	Yes	Yes	Yes
Series 4200vl	Yes	No	No
Series 3500yl	Yes	Yes	Yes
Series 3400cl	Yes	No	No
Series 2800	Yes	No	No
Series 2500	Yes	Yes	Yes
Switch 4000m and 8000m	Yes	Yes	Yes

This chapter describes Traffic/Security filters on the switches covered in this guide. For information on filters for other switches in the above table, refer to the documentation provided for those switches.

Introduction

Feature	Default	Menu	CLI	Web
configure source-port filters	none	n/a	page 12-21	n/a
configure protocol filters	none	n/a	page 12-21	n/a
configure multicast filters	none	n/a	page 12-21	n/a
display filter data	n/a	n/a	page 12-23	n/a

You can enhance in-band security and improve control over access to network resources by configuring static filters to forward (the default action) or drop unwanted traffic. That is, you can configure a traffic filter to either forward or drop all network traffic moving to outbound (destination) ports and trunks (if any) on the switch.

Filter Limits

The switch accepts up to 101 static filters. These limitations apply:

- Source-port filters: up to 78
- Multicast filters: up to 16 with 1024 or fewer VLANs configured. Up to 8 with more than 1024 VLANs configured.
- Protocol filters: up to 7

Using Port Trunks with Filters

The switch manages a port trunk as a single source or destination for source-port filtering. If you configure a port for filtering before adding it to a port trunk, the port retains the filter configuration, but suspends the filtering action while a member of the trunk. If you want a trunk to perform filtering, first configure the trunk, then configure the trunk for filtering. Refer to “Configuring a Filter on a Port Trunk” on page 12-19.

Filter Types and Operation

Table 12-1. Filter Types and Criteria

Static Filter Type	Selection Criteria
Source-Port	Inbound traffic from a designated, physical source-port will be forwarded or dropped on a per-port (destination) basis.
Multicast	Inbound traffic having a specified multicast MAC address will be forwarded to outbound ports (the default) or dropped on a per-port (destination) basis.
Protocol	Inbound traffic having the selected frame (protocol) type will be forwarded or dropped on a per-port (destination) basis.

Source-Port Filters

This filter type enables the switch to forward or drop traffic from *all* end nodes on the indicated source-port to specific destination ports.

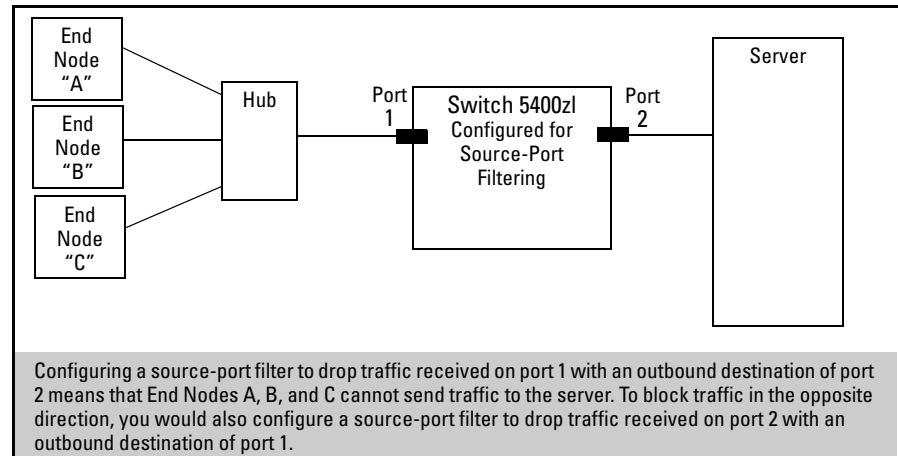


Figure 12-1. Example of a Source-Port Filter Application

Operating Rules for Source-Port Filters

- You can configure one source-port filter for each physical port and port trunk on the switch. (Refer to the **filter** command on page 12-18.)
- You can include all destination ports and trunks in the switch on a single source-port filter.
- Each source-port filter includes:
 - One source port or port trunk (**trk1**, **trk2**, ...**trkn**)
 - A set of destination ports and/or port trunks that includes all untrunked LAN ports and port trunks on the switch
 - An action (forward or drop) for each destination port or port trunk

When you create a source-port filter, the switch automatically sets the filter to forward traffic from the designated source to all destinations for which you do not specifically configure a “drop” action. Thus, it is not necessary to configure a source-port filter for traffic you want the switch to forward unless the filter was previously configured to drop the desired traffic.

- When you create a source port filter, all ports and port trunks (if any) on the switch appear as destinations on the list for that filter, even if routing is disabled and separate VLANs and/or subnets exist. Where traffic would normally be allowed between ports and/or trunks, the switch automatically forwards traffic to the outbound ports and/or trunks you do not specifically configure to drop traffic. (Destination ports that comprise a trunk are listed collectively by the trunk name—such as **Trk1**—instead of by individual port name.)
- Packets allowed for forwarding by a source-port filter are subject to the same operation as inbound packets on a port that is not configured for source-port filtering.
- With multiple IP addresses configured on a VLAN, and routing enabled on the switch, a single port or trunk can be both the source and destination of packets moving between subnets in that same VLAN. In this case, you can prevent the traffic of one subnet from being routed to another subnet of the same port by configuring the port or trunk as both the source and destination for traffic to drop.

Example

If you wanted to prevent server “A” from receiving traffic sent by workstation “X”, but do not want to prevent any other servers or end nodes from receiving traffic from workstation “X”, you would configure a filter to drop traffic from port 5 to port 7. The resulting filter would drop traffic from port 5 to port 7, but would forward all other traffic from any source port to any destination port. (Refer to figures 12-2 and 12-3.




Figure 12-2. 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 12-3. The Filter for the Actions Shown in Figure 12-2

Named Source-Port Filters

You can specify named source-port filters that may be used on multiple ports and port trunks. A port or port trunk can only have one source-port filter, but by using this capability you can define a source-port filter once and apply it to multiple ports and port trunks. This can make it easier to configure and manage source-port filters on your switch. The commands to define, configure, apply, and display the status of named source-port filters are described below.

Operating Rules for Named Source-Port Filters

- A port or port trunk may only have one source-port filter, named or not named.
- A named source-port filter can be applied to multiple ports or port trunks.
- Once a named source-port filter is defined, subsequent changes only modify its action, they don't replace it.
- To change the named source-port filter used on a port or port trunk, the current filter must first be removed, using the **no filter source-port named-filter <filter-name >** command.

- A named source-port filter can only be deleted when it is not applied to any ports.

Defining and Configuring Named Source-Port Filters

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

Syntax: [no] filter source-port named-filter <filter-name>

Defines or deletes a named source-port filter. The <filter-name> may contain a maximum of 20 alpha-numeric characters (longer names may be specified, but they are not displayed.) A filter-name cannot be a valid port or port trunk name.

The maximum number of named source-port filters that can be used is equal to the number of ports on a switch.

*A named source-port filter can only be removed if it is not in use (use the **show filter source-port** command to check the status). Named source-port filters are not automatically deleted when they are no longer used.*

*Use the **no** option to delete an unused named source-port filter.*

Syntax: filter source-port named-filter <filter-name> drop < destination-port-list >

*Configures the named source-port filter to drop traffic having a destination on the ports and/or port trunks in the < destination-port-list >. Can be followed by the **forward** option if you have other destination ports or port trunks previously set to **drop** that you want to change to **forward**. For example:*

filter source-port named-filter <filter-name> drop < 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.*

Using Named Source-Port Filters

A company wants to manage traffic to the Internet and its accounting server on a 26-port switch. Their network is pictured in Figure 12-4. Switch port 1 connects to a router that provides connectivity to a WAN and the Internet. Switch port 7 connects to the accounting server. Two workstations in accounting are connected to switch ports 10 and 11. Two workstations in accounting are connected to switch ports 10 and 11.




Figure 12-4. Network Configuration for Named Source-Port Filters Example

Defining and Configuring Example Named Source-Port Filters. While named source-port filters may be defined and configured in two steps, this is not necessary. Here we define and configure each of the named source-port filters for our example network in a single step.

Traffic/Security Filters and Monitors

Filter Types and Operation

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

ProCurve(config)# show filter source-port

Traffic/Security Filters

Filter Name          | Port List          | Action
-----+-----+-----
web-only            | NOT USED          | drop 2-26
accounting          | NOT USED          | drop 1-6,8-9,12-26
no-incoming-web    | NOT USED          | drop 7,10-11

ProCurve Switch 2626(config)#
```

Ports and port trunks using the filter. When **NOT USED** is displayed the named source-port filter may be deleted.

Lists the ports and port trunks dropped by the filter. Ports and port trunks not shown are forwarded by the filter.

To remove a port or port trunk from the list, update the named source-port filter definition using the **forward** option.

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)#
```

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

```
ProCurve(config)# show filter
```

```
Traffic/Security Filters
```

IDX	Filter Type	Value
1	Source Port	2
2	Source Port	3
3	Source Port	4
4	Source Port	5
5	Source Port	6
6	Source Port	8
7	Source Port	9
8	Source Port	12
20	Source Port	24
21	Source Port	25
22	Source Port	26
23	Source Port	7
24	Source Port	10
25	Source Port	11
26	Source Port	1

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

An automatically assigned index number used to identify the filter for a detailed information listing. A filter retains its assigned IDX number for as long as the filter exists in the switch. The switch assigns the lowest available IDX number to a new filter. This can result in a newer filter having a lower IDX number than an older filter if a previous (source-port or named source-port) filter deletion created a gap in the filter listing.

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.

Traffic/Security Filters and Monitors
Filter Types and Operation

<pre>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 . . .</pre>	<pre>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 . . .</pre>
--	--

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


```
ProCurve(config)# show filter 26
```

Traffic/Security Filters

Filter Type : Source Port
Source Port : 1

Dest	Port Type	Action
1	10/100TX	Forward
2	10/100TX	Forward
3	10/100TX	Forward
4	10/100TX	Forward
5	10/100TX	Forward
6	10/100TX	Forward
7	10/100TX	Drop
8	10/100TX	Forward
9	10/100TX	Forward
10	10/100TX	Drop
11	10/100TX	Drop
12	10/100TX	Forward

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.

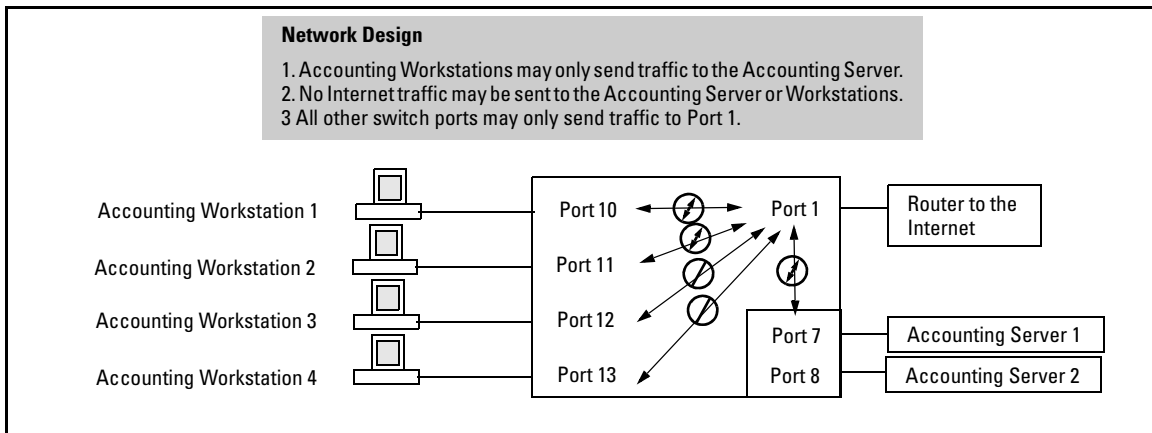


Figure 12-5. Expanded Network Configuration for Named Source-Port Filters Example

Traffic/Security Filters and Monitors

Filter Types and Operation

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)#
```

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)#
```

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)#
```

Static Multicast Filters

This filter type enables the switch to forward or drop multicast traffic to a specific set of destination ports. This helps to preserve bandwidth by reducing multicast traffic on ports where it is unnecessary, and to isolate multicast traffic to enhance security.

You can configure up to 16 static multicast filters (defined by the **filter** command—page 12-21). However, if an IGMP-controlled filter for a joined multicast group has the same multicast address as a static multicast filter configured on a given port, the IGMP-controlled filter overrides the static multicast filter configured on that port. Note that in the default configuration, IGMP is disabled on VLANs configured in the switch. To enable IGMP on a specific VLAN, use the **vlan <vid> ip igmp** command. (For more on this command, refer to the chapter titled “Multimedia Traffic Control with IP Multicast (IGMP)” in the *Multicast and Routing Guide* for your switch.)

The total of static multicast filters and IGMP multicast filters together can range from 389 to 420, depending on the current **max-vlans** setting in the switch. If multiple VLANs are configured, then each filter is counted once per VLAN in which it is used.

Table 12-2. Multicast Filter Limits

Max-VLANs Setting	Maximum # of Multicast Filters (Static and IGMP Combined)
1 (the minimum)	420
8 (the default)	413
32 or higher	389

Notes:

Per-Port IP Multicast Filters. The static multicast filters described in this section filter traffic having a multicast address you specify. To filter all multicast traffic on a per-VLAN basis, refer to the section titled “Configuring and Displaying IGMP” in the chapter titled “Multimedia Traffic Control with IP Multicast (IGMP)” in the *Multicast and Routing Guide* for your switch.

IP Multicast Filters. Multicast filters are configured using the Ethernet format for the multicast address. IP multicast addresses occur in the range of 224.0.0.0 through 239.255.255.255 (which corresponds to the Ethernet multicast address range of 01005e-000000 through 01005e-7fffff). Any static Traffic/Security filters configured with a **multicast** filter type and a multicast address in this range will continue to be in effect unless IGMP learns of a multicast group destination in this range. In this case, IGMP takes over the filtering function for the multicast destination address(es) for as long as the IGMP group is active. If the IGMP group subsequently deactivates, the static filter resumes control over traffic to the multicast address.

Caution

If Spanning Tree is enabled, then the MSTP multicast MAC address (0180c2-000000) should not be filtered. (STP will not operate properly if the MSTP multicast MAC address is filtered.)

Protocol Filters

This filter type enables the switch to forward or drop, on the basis of protocol type, traffic to a specific set of destination ports on the switch. Filtered protocol types include:

- AppleTalk
- ARP
- IPX
- NetBEUI
- SNA

Only one filter for a particular protocol type can be configured at any one time. For example, a separate protocol filter can be configured for each of the protocol types listed above, but only one of those can be an IP filter. Also, the destination ports for a protocol filter can be on different VLANs.

You can configure up to seven protocol filters.

Configuring Traffic/Security Filters

Use this procedure to specify the type of filters to use on the switch and whether to forward or drop filtered packets for each filter you specify.

1. Select the static filter type(s).
2. For inbound traffic matching the filter type, determine the filter action you want for each outbound (destination) port on the switch (forward or drop). The default action for a new filter is to forward traffic of the specified type to all outbound ports.
3. Configure the filter.
4. Use **show filter** (page 12-23) to check the filter listing to verify that you have configured correct action for the desired outbound ports.

Configuring a Source-Port Traffic Filter

Syntax: [no] filter

[source-port < port-number | trunk-name>]

*Specifies one inbound port or trunk. Traffic received inbound on this interface from other devices will be filtered. The **no** form of the command deletes the source-port filter for < port-number > and returns the destination ports for that filter to the **Forward** action. (Default: Forward on all ports.)*

Note: *If multiple VLANs are configured, the source-port and the destination port(s) must be in the same VLAN unless routing is enabled. Similarly, if a VLAN containing both the source and destination is multi-netted, the source and destination ports and/or trunks must be in the same subnet unless routing is enabled.*

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

*Configures the filter to drop traffic for the ports and/or trunks in the designated < destination-port-list >. Can be followed by **forward** < destination-port-list > if you have other destination ports set to **drop** that you want to change to **forward**. If no drop or forward action is specified, the switch automatically creates a filter with a **forward** action from the designated source port (or trunk) to all destination ports (or trunks) on the switch.*

[forward] < port-list >

*Configures the filter to forward traffic for the ports and/or trunks in the designated < destination-port-list >. Because **forward** is the default state for destinations in a filter, this command is useful when destinations in an existing filter are configured for **drop** and you want to change them to **forward**. Can be followed by **drop** < destination-port-list > if you have other destination ports set to **forward** that you want to change to **drop**. If no drop or forward action is specified, the switch automatically creates a filter with a forward action from the designated source port (or trunk) to all destination ports (or trunks) on the switch.*

Example of Creating a Source-Port Filter

For example, assume that you want to create a source-port filter that drops all traffic received on port 5 with a destination of port trunk 1 (**Trk1**) and any port in the range of port 10 to port 15. To create this filter you would execute this command:

```
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 is used for configuring a filter on an individual port. However, the configuration process requires two steps:

1. Configure the port trunk.
2. Configure a filter on the port trunk by using the trunk name (**trk1**, **trk2**, ...**trk6**) instead of a port name.

For example, to create a filter on port trunk 1 to drop traffic received inbound for trunk 2 and ports 10-15:

```
ProCurve(config)# filter source-port trk1 drop trk2,10-15
```

Note that if you first configure a filter on a port and then later add the port to a trunk, the port remains configured for filtering *but the filtering action will be suspended while the port is a member of the trunk*. That is, the trunk does not adopt filtering from the port configuration. You must still explicitly configure the filter on the port trunk. If you use the **show filter < index >** command for a filter created before the related source port was added to a trunk, the port number appears between asterisks (*), indicating that the filter action has been suspended for that filter. For example, if you create a

filter on port 5, then create a trunk with ports 5 and 6, and display the results, you would see the following:

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

Traffic/Security Filters

IDX	Filter Type	Value
1	Source Port	*5*

```
ProCurve(config)# show filter 1
```

Traffic/Security Filters

Filter Type : Source Port
Source Port : *5*

Dest Port	Type	Action
1	100/1000T	Forward
2	100/1000T	Drop
3	100/1000T	Forward
4	100/1000T	Forward
.	.	.
.	.	.
.	.	.

The *5* shows that port 5 is configured for filtering, but the filtering action has been suspended while the port is a member of a trunk. If you want the trunk to which port 5 belongs to filter traffic, then you must explicitly configure filtering on the trunk.

Note: If you configure an existing trunk for filtering and later add another port to the trunk, the switch will apply the filter to all traffic moving on any link in the trunk. If you remove a port from the trunk it returns to the configuration it had before it was added to the trunk

Figure 12-6. Example of Switch Response to Adding a Filtered Source Port to a Trunk

Editing a Source-Port Filter

The switch includes in one filter the action(s) for all destination ports and/or trunks configured for a given source port or trunk. Thus, if a source-port filter already exists and you want to change the currently configured action for some destination ports or trunks, use the **filter source-port** command to update the existing filter. For example, suppose you configure a filter to drop traffic received on port 8 and destined for ports 1 and 2. The resulting filter is shown on the left in figure 12-7. 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 12-7:

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 12-7. Assigning Additional Destination Ports to an Existing Filter

Configuring a Multicast or Protocol Traffic Filter

Syntax: [no] filter

[multicast < mac-address >]

Specifies a multicast address. Inbound traffic received (on any port) with this multicast address will be filtered. (Default: Forward on all ports.)

*The **no** form of the command deletes the multicast filter for the < mac-address > multicast address and returns the destination ports for that filter to the **Forward** action.*

[< forward | drop > < port-list >]

Specifies whether the designated destination port(s) should forward or drop the filtered traffic.

[protocol < ip | ipx | arp | appletalk | sna | netbeui >]

Specifies a protocol type. Traffic received (on any port) with this protocol type will be filtered. (Default: Forward on all ports.)

*The **no** form of the command deletes the protocol filter for the specified protocol and returns the destination ports for that filter to the **Forward** action.*

[< forward | drop > < port-list >]

Specifies whether the designated destination port(s) should forward or drop the filtered traffic.

For example, suppose you wanted to configure the filters in table 12-3 on a switch. (For more on source-port filters, refer to “Configuring a Source-Port Traffic Filter” on page 12-18.)

Table 12-3. Filter Example

Filter Type	Filter Value	Action	Destination Ports
Source-Port	Inbound ports: A1, A2*	Drop	D1-D4
Multicast	010000-123456	Drop	C1-C24, D5-D10
Multicast	010000-224466	Drop	B1-B4
Protocol	Appletalk	Drop	C12-C18, D1
Protocol	ARP	Drop	D17, D21-D24

*Because the switch allows one inbound port in a source-port filter, the requirement to filter ports A1 and A2 means you will configure two separate source-port filters.

The following commands configure the filters listed above:

```
ProCurve(config)# filter source-port a1 drop e d1-d4
ProCurve(config)# filter source-port a2 drop d1-d4
ProCurve(config)# filter multicast 010000-123456 drop e c1-c24,d5-d10
ProCurve(config)# filter multicast 010000-224466 drop e b1-b4
ProCurve(config)# filter protocol appletalk drop e c12-c18,d1
ProCurve(config)# filter protocol arp drop e d17,d21-d24
```

Figure 12-8. Configuring Various Traffic/Security Filters

Filter Indexing

The switch automatically assigns each new filter to the lowest-available index (IDX) number. The index numbers are included in the **show filter** command described in the next section and are used with the **show filter < index >** command to display detailed information about a specific filter.

If there are no filters currently configured, and you create three filters in succession, they will have index numbers 1 - 3. However, if you then delete the filter using index number “2” and then configure two new filters, the first new filter will receive the index number “2” and the second new filter will receive the index number “4”. This is because the index number “2” was made vacant by the earlier deletion, and was therefore the lowest index number available for the next new filter.

Displaying Traffic/Security Filters

This command displays a listing of all filters by index number and also enables you to use the index number to display the details of individual filters.

Syntax: show filter

Lists the filters configured in the switch, with corresponding filter index (IDX) numbers.

IDX: *An automatically assigned index number used to identify the filter for a detailed information listing. A filter retains its assigned IDX number for as long as the filter exists in the switch. The switch assigns the lowest available IDX number to a new filter. This can result in a newer filter having a lower IDX number than an older filter if a previous filter deletion created a gap in the filter listing.*

Filter Type: *Indicates the type of filter assigned to the IDX number (source-port, multicast, or protocol).*

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

[*index*]

Lists the filter type and other data for the filter corresponding to the index number in the **show filter** output. Also lists, for each outbound destination port in the switch, the port number, port type, and filter action (forward or drop). The switch assigns the lowest available index number to a new filter. If you delete a filter, the index number for that filter becomes available for the next filter you create.

For example, to display the filters created in figure 12-8 on page 12-22 and then list the details of the multicast filter for multicast address **010000-224466**:

Traffic/Security Filters and Monitors
 Configuring Traffic/Security Filters

```

ProCurve(config)# show filter
Traffic/Security Filters
  (IDX) Filter Type | Value
-----+-----
  1 Source Port | A1
  2 Source Port | A2
  3 Multicast | 010000-123456
  4 Multicast | 010000-224466
  5 Protocol | AppleTalk
  6 Protocol | ARP
    
```

Filter Index Numbers (Automatically Assigned)

Criteria for Individual Filters

```

ProCurve(config)# show filter 4
Traffic/Security Filters
Filter Type : Multicast
Multi-cast Address : 010000-224466
    
```

Uses the index number (IDX) for a specific filter to list the details for that filter only.

Dest Port	Type	Action
A1	1000LX	Forward
A2		Forward
A3		Forward
A4	1000SX	Forward
B1	100/1000T	Drop
B2	100/1000T	Drop
B3	100/1000T	Drop
B4	100/1000T	Drop
C1	10/100TX	Forward
C2	10/100TX	Forward
C3	10/100TX	Forward
C4	10/100TX	Forward
C5	10/100TX	Forward
C6	10/100TX	Forward
C7	10/100TX	Forward

-- MORE --, next page: Space, next line: Enter.

Figure 12-9. Example of Displaying Filter Data

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

Contents

Overview	13-3
Why Use Port-Based or User-Based Access Control?	13-3
General Features	13-3
User Authentication Methods	13-4
802.1X User-Based Access Control	13-4
802.1X Port-Based Access Control	13-5
Alternative To Using a RADIUS Server	13-6
Accounting	13-6
Terminology	13-6
General 802.1X Authenticator Operation	13-9
Example of the Authentication Process	13-9
VLAN Membership Priority	13-10
General Operating Rules and Notes	13-12
General Setup Procedure for 802.1X Access Control	13-14
Do These Steps Before You Configure 802.1X Operation	13-14
Overview: Configuring 802.1X Authentication on the Switch	13-15
Configuring Switch Ports as 802.1X Authenticators	13-16
1. Enable 802.1X Authentication on Selected Ports	13-17
A. Enable the Selected Ports as Authenticators and Enable the (Default) Port-Based Authentication	13-17
B. Specify User-Based Authentication or Return to Port-Based Authentication	13-18
Example: Configuring User-Based 802.1X Authentication	13-19
Example: Configuring Port-Based 802.1X Authentication	13-19
2. Reconfigure Settings for Port-Access	13-19

3. Configure the 802.1X Authentication Method	13-21
4. Enter the RADIUS Host IP Address(es)	13-22
5. Enable 802.1X Authentication on the Switch	13-23
6. Optional: Reset Authenticator Operation	13-23
7. Optional: Configure 802.1X Controlled Directions	13-24
Wake-on-LAN Traffic	13-24
Operating Notes	13-25
Example: Configuring 802.1X Controlled Directions	13-25
802.1X Open VLAN Mode	13-26
Introduction	13-26
VLAN Membership Priorities	13-27
Use Models for 802.1X Open VLAN Modes	13-28
Operating Rules for Authorized-Client and Unauthorized-Client VLANs	13-33
Setting Up and Configuring 802.1X Open VLAN Mode	13-37
802.1X Open VLAN Operating Notes	13-41
Option For Authenticator Ports: Configure Port-Security To Allow Only 802.1X-Authenticated Devices	13-42
Port-Security	13-43
Configuring Switch Ports To Operate As Supplicants for 802.1X Connections to Other Switches	13-44
Example	13-44
Supplicant Port Configuration	13-46
Displaying 802.1X Configuration, Statistics, and Counters	13-48
Show Commands for Port-Access Authenticator	13-48
Viewing 802.1X Open VLAN Mode Status	13-51
Show Commands for Port-Access Supplicant	13-55
How RADIUS/802.1X Authentication Affects VLAN Operation .	13-56
Operating Notes	13-60
Messages Related to 802.1X Operation	13-61

Overview

Feature	Default	Menu	CLI	Web
Configuring Switch Ports as 802.1X Authenticators	Disabled	n/a	page 13-16	n/a
Configuring 802.1X Open VLAN Mode	Disabled	n/a	page 13-26	n/a
Configuring Switch Ports to Operate as 802.1X Supplicants	Disabled	n/a	page 13-44	n/a
Displaying 802.1X Configuration, Statistics, and Counters	n/a	n/a	page 13-48	n/a
How 802.1X Affects VLAN Operation	n/a	n/a	page 13-56	n/a
RADIUS Authentication and Accounting	Refer to chapter 6, "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 32 authenticated clients per-port.

- Port-Based access control option allowing authentication by a single client to open the port. This option does not force a client limit and, on a port opened by an authenticated client, allows unlimited client access without requiring further authentication.
- Supplicant implementation using CHAP authentication and independent user credentials on each port.
- Local authentication of 802.1X clients using the switch's local username and password (as an alternative to RADIUS authentication).
- 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 32 802.1X-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 32 per port), where each client gains access to the LAN by entering valid user credentials. This operation improves security by opening a given port only to individually authenticated clients, while simultaneously blocking access to the same port for clients that cannot be authenticated. All sessions must use the same untagged VLAN. Also, an authenticated client can use any tagged VLAN memberships statically configured on the port, provided the client is configured to use the tagged VLAN memberships available on the port. (Note that

the session total includes any sessions begun by the Web Authentication or MAC Authentication features covered in chapter 4.) For more information, refer to “Option For Authenticator Ports: Configure Port-Security To Allow Only 802.1X-Authenticated Devices” on page 13-42.

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 13-4) should be used instead of port-based access control. Using the user-based method enables you to specify up to 32 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 Notes” on page 13-60.

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.

Accounting

The switches covered in this guide also provide RADIUS Network accounting for 802.1X access. Refer to chapter 6, “RADIUS-Administered CoS and Rate-Limiting”.

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 13-26.)

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 13-4.

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 13-1 on page 13-11.

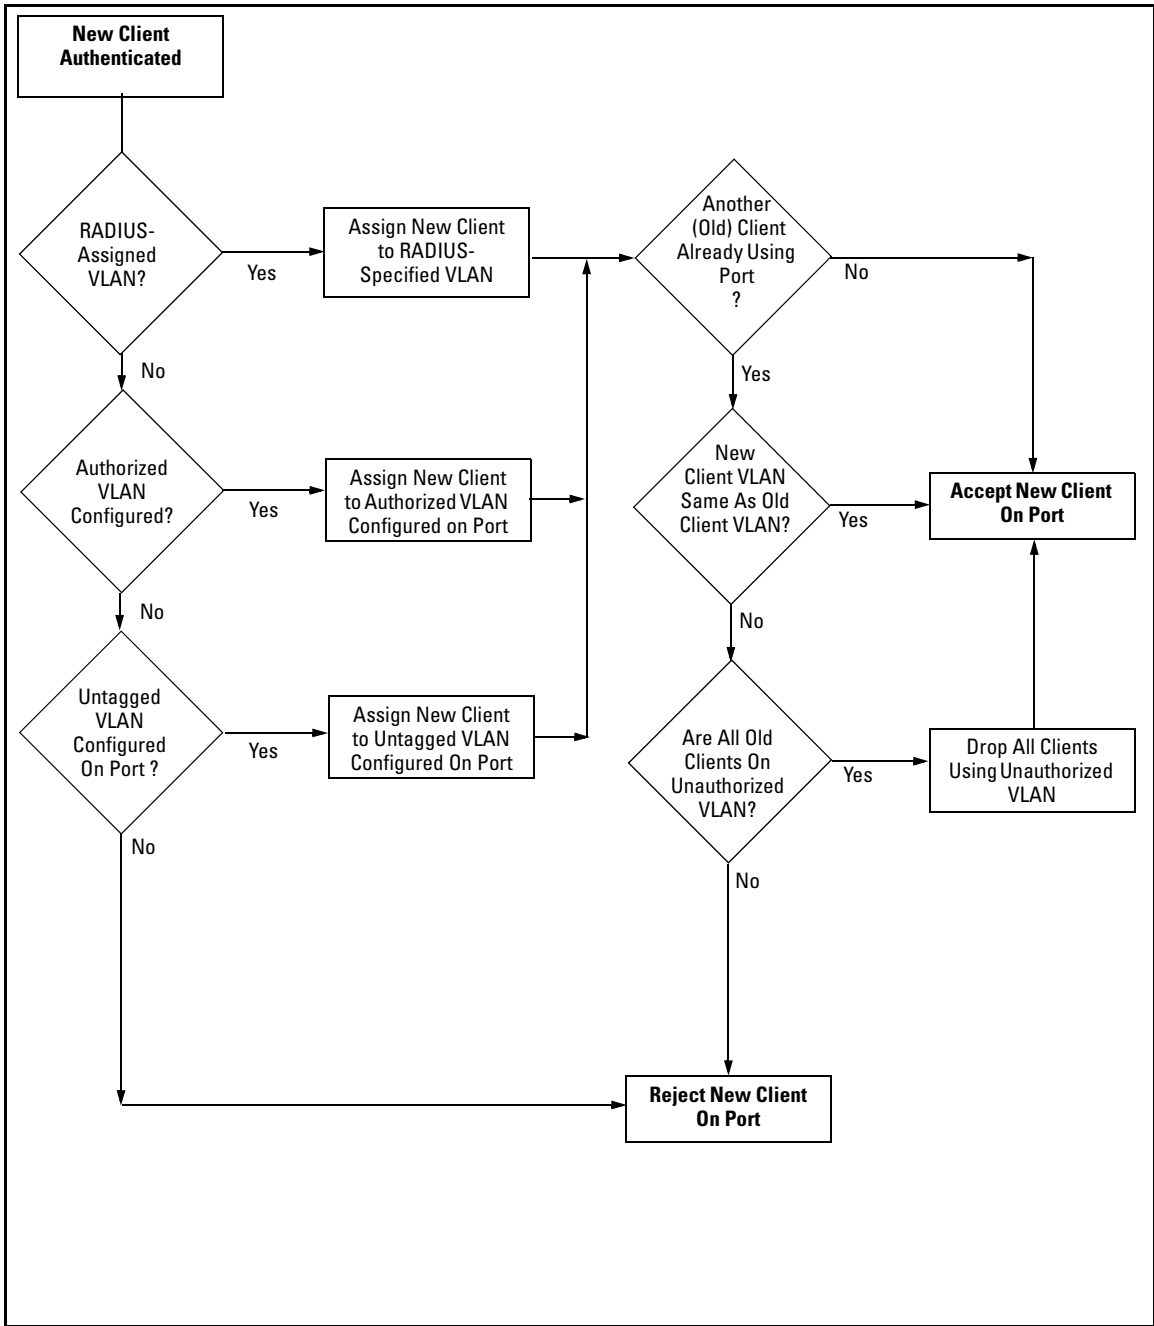


Figure 13-1. Priority of VLAN Assignment for an Authenticated Client

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 13-26.)

- 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 13-26. (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 13-44.
- To help maintain security, 802.1X and LACP cannot both be enabled on the same port. If you try to configure 802.1X on a port already configured for LACP (or the reverse) you will see a message similar to the following:

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

General Setup Procedure for 802.1X Access Control

Do These Steps Before You Configure 802.1X Operation

1. Configure a local username and password on the switch for both the Operator (login) and Manager (enable) access levels. (While this may or may not be required for your 802.1X configuration, HP recommends that you use a local username and password pair at least until your other security measures are in place.)
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 13-17.)
3. Determine whether to use user-based access control (page 13-4) or port-based access control (page 13-5).
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 13-26.
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 13-4
 - “802.1X Port-Based Access Control” on page 13-5
 - “Configuring Switch Ports To Operate As Supplicants for 802.1X Connections to Other Switches” on page 13-44.
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 13-17.
 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 13-26.
 3. Configure the 802.1X authentication type. Options include:
 - Local Operator username and password (the default). This option allows a client to use the switch’s local username and password as valid 802.1X credentials for network access.
 - EAP RADIUS: This option requires your RADIUS server application to support EAP authentication for 802.1X.
 - CHAP (MD5) RADIUS: This option requires your RADIUS server application to support CHAP (MD5) authentication.Refer to page 13-21.
 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 13-22.
 5. Enable 802.1X authentication on the switch. Refer to “1. Enable 802.1X Authentication on Selected Ports” on page 13-17.
 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 13-42.
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 13-44.)

Configuring Switch Ports as 802.1X Authenticators

802.1X Authentication Commands	Page
[no] aaa port-access authenticator < <i>port-list</i> >	13-17
[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]	13-17
aaa authentication port-access < local eap-radius chap-radius >	13-21
[no] aaa port-access authenticator active	13-16
aaa port-access < <i>port-list</i> > controlled-directions <both in>	13-24
[no] port-security [ethernet] < <i>port-list</i> > learn-mode port-access	13-42
802.1X Open VLAN Mode Commands	13-26
802.1X Supplicant Commands	13-44
802.1X-Related Show Commands	13-48
RADIUS server configuration	13-22

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 13-23 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 13-23.*

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 - 32 >`

*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 13-17.) 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 13-2. 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 a13-a15
ProCurve(config)# no aaa port-access authenticator a13-a15 client-limit
```

Figure 13-3. 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 download this software and begin the authentication process. Refer to "802.1X Open VLAN Mode" on page 13-26.)

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)

— Continued —

aaa port-access authenticator < port-list >

[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 13-26.

aaa port-access authenticator < port-list >

[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 13-26.

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.*

—Continued—

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.)

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

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

Status and Counters - Authentication Information

Login Attempts : 3
 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			

Annotations:
 - Arrow from top right: Configuration command for EAP-RADIUS authentication.
 - Arrow from middle right: 802.1X (Port-Access) configured for EAP-RADIUS authentication.

Figure 13-4. 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 6, “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 4, “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 13-8.
- 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 13-5. Example of Configuring 802.1X Controlled Directions

802.1X Open VLAN Mode

802.1X Authentication Commands	page 13-16
802.1X Supplicant Commands	page 13-46
802.1X Open VLAN Mode Commands	
[no] aaa port-access authenticator < <i>port-list</i> >	page 13-40
[auth-vid < <i>vlan-id</i> >]	
[unauth-vid < <i>vlan-id</i> >]	
802.1X-Related Show Commands	page 13-48
RADIUS server configuration	pages 13-22

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 13-2. 802.1X Open VLAN Mode Options

802.1X Per-Port Configuration	Port Response
No Open VLAN mode:	The port automatically blocks a client that cannot initiate an authentication session.

Open VLAN mode with **both** of the following configured:

- Unauthorized-Client VLAN • When the port detects a client 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.

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.

802.1X Per-Port Configuration	Port Response
Authorized-Client VLAN	<ul style="list-style-type: none"><li data-bbox="565 210 1272 288">• After client authentication, the port drops membership in the Unauthorized-Client VLAN and becomes an untagged member of this VLAN. <p data-bbox="594 305 1272 517">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 13-21—to delay moving the port into the Unauthorized-Client VLAN.</p> <p data-bbox="594 552 1272 656">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"><li data-bbox="565 673 1272 777">• 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.<li data-bbox="565 795 1272 1024">• 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:

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

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

- If the port is statically configured as a tagged member of a VLAN, the port returns to tagged membership in this VLAN upon successful client authentication. This happens even if the RADIUS server assigns the port to another, authorized VLAN. Note that if the port is already configured as a tagged member of a VLAN that RADIUS assigns as an authorized VLAN, then the port becomes an *untagged* member of that VLAN for the duration of the client connection.

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 13-1 on page 13-11.

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

802.1X Open VLAN Mode

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.	
<p>Note: An authorized-client VLAN configuration can be overridden by a RADIUS authentication that assigns a VLAN. (Refer to figure 13-1 on page 13-11.)</p>	

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 13-28.)
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.

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

802.1X Open VLAN Mode

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 Unauthenticated-Client VLAN and joins the appropriate VLAN. (Refer to “VLAN Membership Priorities” on page 13-27.• 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 13-28.) <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>

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	When a new client is authenticated on a given port: <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.)

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

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 32 clients per-port. The Unauthorized-Client VLAN feature can operate on an 802.1X-configured port regardless of how many clients the port is configured to support. However, all clients on the same port must operate through the same untagged VLAN membership. This means that any client accessing a given port must be able to authenticate and operate on the same VLAN as any other previously authenticated clients that are currently using the port. Thus, an Unauthorized-Client VLAN configured on a switch port that allows multiple 802.1X clients cannot be used if there is already an authenticated client using the port on another VLAN. Also, a client using the Unauthenticated-Client VLAN will be blocked when another client becomes authenticated on the port. For this reason, the best utilization of the Unauthorized-Client VLAN feature is in instances where only one client is allowed per-port. Otherwise, unauthenticated clients are subject to being blocked at any time by authenticated clients using a different VLAN. (Using the same VLAN for authenticated and unauthenticated clients can create a security risk and is not recommended.)

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 13-2 on page 13-29 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 13-17.) 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 13-42.

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 13-37.

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 13-51.

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
```

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

Control mode required for Port-Security Support

Figure 13-6. 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 14), 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 - 32 >`

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 13-16.)

— 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 13-4.)

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

802.1X Authentication Commands	page 13-16
802.1X Supplicant Commands	
[no] aaa port-access < supplicant < [ethernet] < port-list >	page 13-46
[auth-timeout held-period start-period max-start initialize identity secret clear-statistics]	page 13-46
802.1X-Related Show Commands	page 13-48
RADIUS server configuration	pages 13-22

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 13-7. 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 13-44 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 13-16
802.1X Supplicant Commands	page 13-44
802.1X Open VLAN Mode Commands	page 13-26
802.1X-Related Show Commands	
show port-access authenticator	below
show port-access supplicant	page 13-55
Details of 802.1X Mode Status Listings	page 13-51
RADIUS server configuration	pages 13-22

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. Includes the port traffic priority (CoS) assigned to inbound traffic and the **rate-limit** settings, if any, specified by a RADIUS server for a current 802.1X authenticated client session. (Refer to “RADIUS-Administered CoS and Rate-Limiting” on page 6-4 in this guide.)*
- *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 13-19. 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 13-46.) 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 13-19)*
- *Unauth-VLAN ID (if any)*
- *Auth-VLAN ID (if any)*

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

```

ProCurve(config)# show port-access authenticator config

Port Access Authenticator Configuration

Port-access authenticator activated [No] : No

      | Re-auth Access  Max  Quiet  TX      Supplicant  Server  Cntrl
Port | Period Control Reqs  Period Timeout Timeout  Timeout  Dir
----+-----
1   | No    Auto    2    60    30     30     30     both
2   | No    Auto    2    60    30     30     30     in
    
```

Figure 13-8. Example of show port-access authenticator config Command

Table 13-3. Field Descriptions of show port-access authenticator config Command Output (Figure 13-9)

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.
Supplicant Timeout	Period of time (in seconds) that the switch waits for a supplicant response to an EAP request.
Server Timeout	Period of time (in seconds) that the switch waits for a server response to an authentication request.
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 as illustrated in figure 13-9. Table 13-2 describes the data that these two commands display. Figure 13-10 shows related VLAN data that can help you to see how the switch is using statically configured VLANs to support 802.1X operation.

```

ProCurve 3400(config)# show port-access authenticator vlan
Port Access Authenticator VLAN Configuration

Port-access authenticator activated [No] : Yes

   Access   Unauth   Auth
Port Control VLAN ID  VLAN ID
-----
1   Auto   (100)   101
2   Auto   100     (101)
3   Auto   100     (0)
4   Auto   100     101
    
```

In these two **show** outputs, an Unauth VLAN ID appearing in the Current VLAN ID column for the same port indicates an unauthenticated client is connected to this port. (Assumes that the port is not a statically configured member of VLAN 100.)

```

HP3400# show port-access authenticator 1-4
Port Access Authenticator Status

Port-access authenticator activated [No] : No

Port Status  Authenticator  Authenticator  Current  Current  % Curr. Rate
-----
1   Closed Connecting  Idle           (100)   No-override No-override
2   (Open) (Authorized)  Idle           (101)   No-override No-override
3   Closed Connecting  Idle           100     No-override No-override
4   Closed Disconnected Idle           (No PVID)  No-override No-override
    
```

Note: Series 5400zl switches do not include the **Authenticator State** and **Authenticator Backend State** fields shown in this figure.

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

1. **Open** in the Status column
2. **Authorized** in the Authenticator State column
3. The Auth VLAN ID (**101**) is also in the Current VLAN ID column. (This assumes that the port is not a statically configured member of VLAN 101.)
4. A "0" in the row for port 3 indicates there is no Authorized VLAN configured for port 3.
5. No PVID" means there is currently no untagged VLAN membership on port 4.

Figure 13-9. Example Showing Ports Configured for Open VLAN Mode

Thus, in the output shown in figure 13-9:

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

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

Table 13-1. Output for Determining Open VLAN Mode Status (Figure 13-9, Upper)

Status Indicator	Meaning
Access Control	
This state is controlled by the following port-access command syntax:	
ProCurve(config)# aaa port-access authenticator < port-list > control < authorized auto unauthorized >	
Auto: Configures the port to allow network access to any connected device that supports 802.1X authentication and provides valid 802.1X credentials. (This is the default authenticator setting.)	
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.)	
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.	
Unauthorized VLAN ID	< vlan-id >: Lists the VID of the static VLAN configured as the unauthorized VLAN for the indicated port. 0: No unauthorized VLAN has been configured for the indicated port.
Authorized VLAN ID	< vlan-id >: Lists the VID of the static VLAN configured as the authorized VLAN for the indicated port. 0: No authorized VLAN has been configured for the indicated port.

Table 13-3. Output for Determining Open VLAN Mode Status (Figure 13-9, Lower)

Status Indicator	Meaning
Status	Closed: Either no client is connected or the connected client has not received authorization through 802.1X authentication. Open: An authorized 802.1X supplicant is connected to the port.
Current VLAN ID	< vlan-id >: Lists the VID of the static, untagged VLAN to which the port currently belongs. No PVID: The port is not an untagged member of any VLAN.
Current Port CoS	<i>Refer to the section describing RADIUS support for Identity-Driven Management—IDM—in chapter 6, “RADIUS Authentication and Accounting” in this guide.</i>
% Curr. Rate Limit Inbound	

Syntax: show vlan < vlan-id >

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

```

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

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

Overridden Port VLAN configuration

Port Mode
-----
B1  Untagged
B3  Untagged
    
```

Note that ports B1 and B3 are not in the upper listing, but are included under "Overridden Port VLAN configuration". This shows that static, untagged VLAN memberships on ports B1 and B3 have been overridden by temporary assignment to the authorized or unauthorized VLAN. Using the **show port-access authenticator < port-list >** command shown in figure 13-9 provides details.

Figure 13-10. 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 13-46).

For descriptions of the supplicant parameters, refer to “Configuring a Supplicant Switch Port” on page 13-46.

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.)

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

```

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

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

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

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

Figure 13-11. Example of an Active VLAN Configuration

In figure 13-11, 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).

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

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

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

```

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

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

Overridden Port VLAN configuration

Port  Mode
----  ---
A2    No
  
```

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

Note: With the current VLAN configuration (figure 13-11), the only time port A2 appears in this **show vlan 22** listing is during an 802.1X session with an attached client. Otherwise, port A2 is not listed.

Figure 13-12. The Active Configuration for VLAN 22 Temporarily Changes for the 802.1X Session

- With the preceding in mind, since (static) VLAN 33 is configured as untagged on port A2 (see figure 13-11), and since a port can be untagged on only one VLAN, port A2 loses access to VLAN 33 for the duration of the 802.1X session involving VLAN 22. You can verify the temporary loss of access to VLAN 33 with the **show vlan 33** command.

Even though port A2 is configured as Untagged on (static) VLAN 33 (see figure 13-11), it does not appear in the VLAN 33 listing while the 802.1X session is using VLAN 22 in the Untagged status. However, after the 802.1X session with VLAN 22 ends, the active configuration returns port A2 to VLAN 33.

```

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

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

Overridden Port VLAN configuration

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

Figure 13-13. The Active Configuration for VLAN 33 Temporarily Drops Port 22 for the 802.1X Session

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

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

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

After the 802.1X session on VLAN 22 ends, the active configuration again includes VLAN 33 on port A2.

Figure 13-14. The Active Configuration for VLAN 33 Restores Port A2 After the 802.1X Session Ends

Notes

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

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

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

Operating Notes

- **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 13-4. 802.1X Operating Messages

Message	Meaning
Port < <i>port-list</i> > is not an authenticator.	The ports in the port list have not been enabled as 802.1X authenticators. Use this command to enable the ports as authenticators: <pre>ProCurve (config) # aaa port-access authenticator e 10</pre>
Port < <i>port-list</i> > is not a supplicant.	Occurs when there is an attempt to change the supplicant configuration on a port that is not currently enabled as a supplicant. Enable the port as a supplicant and then make the desired supplicant configuration changes. Refer to “Enabling a Switch Port as a Supplicant” on page 13-46.
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 6-47).
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 < <i>port-number</i> >: 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.

— This page is intentionally unused —

Configuring and Monitoring Port Security

Contents

Overview	14-3
Port Security	14-4
Basic Operation	14-4
Eavesdrop Protection	14-5
Blocking Unauthorized Traffic	14-5
Trunk Group Exclusion	14-6
Planning Port Security	14-7
Port Security Command Options and Operation	14-8
Port Security Display Options	14-8
Configuring Port Security	14-12
Retention of Static Addresses	14-18
MAC Lockdown	14-23
Differences Between MAC Lockdown and Port Security	14-25
MAC Lockdown Operating Notes	14-26
Deploying MAC Lockdown	14-27
MAC Lockout	14-31
Port Security and MAC Lockout	14-33
Web: Displaying and Configuring Port Security Features	14-34
Reading Intrusion Alerts and Resetting Alert Flags	14-34
Notice of Security Violations	14-34
How the Intrusion Log Operates	14-35
Keeping the Intrusion Log Current by Resetting Alert Flags	14-36
Menu: Checking for Intrusions, Listing Intrusion Alerts, and Resetting Alert Flags	14-37
CLI: Checking for Intrusions, Listing Intrusion Alerts, and Resetting Alert Flags	14-38
Using the Event Log To Find Intrusion Alerts	14-40

Configuring and Monitoring Port Security
Contents

Web: Checking for Intrusions, Listing Intrusion Alerts, and Resetting Alert Flags 14-41

Operating Notes for Port Security 14-42

Overview

Feature	Default	Menu	CLI	Web
Displaying Current Port Security	n/a	—	page 14-8	page 14-34
Configuring Port Security	disabled	—	page 14-12	page 14-34
Retention of Static Addresses	n/a	—	page 14-18	n/a
MAC Lockdown	disabled	—	page 14-23	
MAC Lockout	disabled	—	page 14-31	
Intrusion Alerts and Alert Flags	n/a	page 14-40	page 14-38	page 14-41

Port Security (Page 14-4). This feature enables you to configure each switch port with a unique list of the MAC addresses of devices that are authorized to access the network through that port. This enables individual ports to detect, prevent, and log attempts by unauthorized devices to communicate through the switch.

Note

This feature does not prevent intruders from receiving broadcast and multi-cast traffic. Also, Port Security and MAC Lockdown are mutually exclusive on a switch. If one is enabled, then the other cannot be used.

MAC Lockdown (Page 14-23). This feature, also known as “Static Addressing”, is used to prevent station movement and MAC address “hijacking” by allowing a given MAC address to use only an assigned port on the switch. MAC Lockdown also restricts the client device to a specific VLAN. (See also the **Note**, above.)

MAC Lockout (Page 14-31). This feature enables you to block a specific MAC address so that the switch drops all traffic to or from the specified address.

Port Security

Basic Operation

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

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

Eavesdrop Protection. Using either the port-security command or the switch’s web browser interface to enable port security on a given port automatically enables eavesdrop prevention on that port.

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

- Alert flags that are captured by network management tools such as ProCurve Manager (PCM and PCM+)
- Alert Log entries in the switch’s web browser interface
- Event Log entries in the console interface
- Intrusion Log entries in the menu interface, CLI, or web browser interface

For any port, you can configure the following:

- **Action:** Used when a port detects an intruder. Specifies whether to send an SNMP trap to a network management station and whether to disable the port.
- **Address Limit:** Sets the number of authorized MAC addresses allowed on the port.
- **Learn-Mode:** Specify how the port acquires authorized addresses.
 - **Continuous:** Allows the port to learn addresses from inbound traffic from any connected device. This is the default setting.
 - **Limited-Continuous:** Sets a finite limit (1 - 32) to the number of learned addresses allowed per port.

- **Static:** Enables you to set a fixed limit on the number of MAC addresses authorized for the port and to specify some or all of the authorized addresses. (If you specify only some of the authorized addresses, the port learns the remaining authorized addresses from the traffic it receives from connected devices.)
- **Configured:** Requires that you specify all MAC addresses authorized for the port. The port is not allowed to learn addresses from inbound traffic.
- **Authorized (MAC) Addresses:** Specify up to eight devices (MAC addresses) that are allowed to send inbound traffic through the port. This feature:
 - Closes the port to inbound traffic from any unauthorized devices that are connected to the port.
 - Provides the option for sending an SNMP trap notifying of an attempted security violation to a network management station and, optionally, disables the port. (For more on configuring the switch for SNMP management, see “Trap Receivers and Authentication Traps” in the *Management and Configuration Guide* for your switch.)
- **Port Access:** Allows only the MAC address of a device authenticated through the switch’s 802.1X Port-Based access control. Refer to chapter 13, Configuring Port-Based and User-Based Access Control (802.1X).

For configuration details, refer to “Configuring Port Security” on page 14-12.

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:

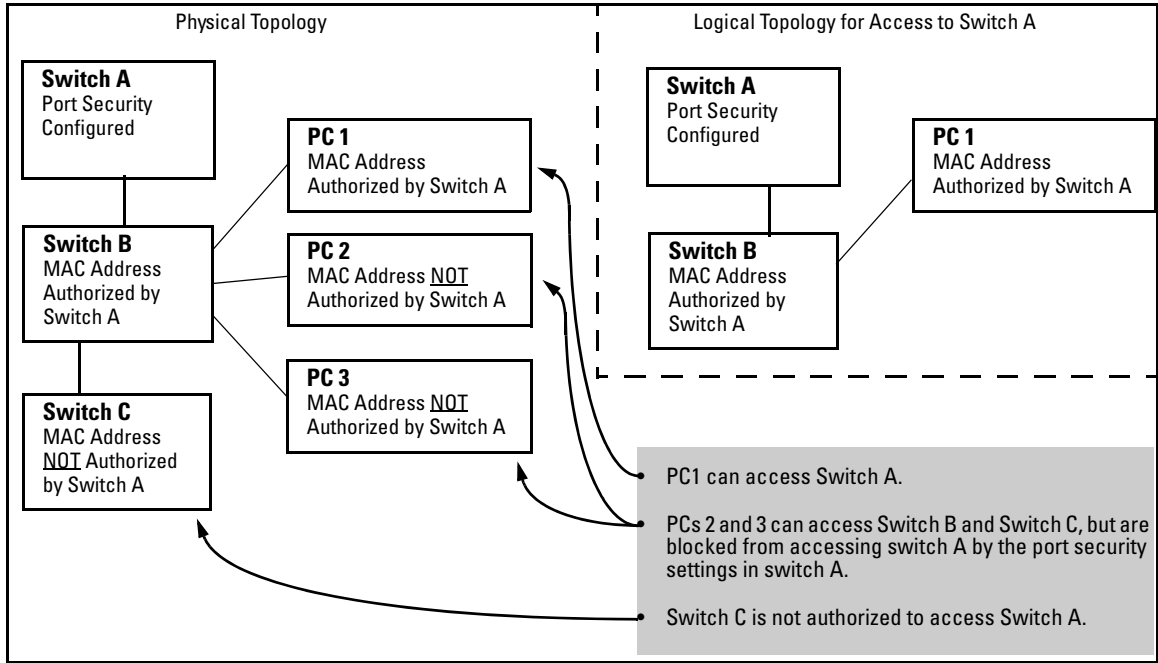


Figure 14-1. Example of How Port Security Controls Access

Note

Broadcast and Multicast traffic is always allowed, and can be read by intruders connected to a port on which you have configured port security.

Trunk Group Exclusion

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

Planning Port Security

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

Port Security Command Options and Operation

Port Security Commands Used in This Section

show port-security	14-9
show mac-address	
port-security	14-12
< <i>port-list</i> >	14-12
learn-mode	14-12
address-limit	14-15
mac-address	14-16
action	14-16
clear-intrusion-flag	14-17
no port-security	14-17

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.

Port Security Display Options

You can use the CLI to display the current port-security settings and to list the currently authorized MAC addresses the switch detects on one or more ports.

Displaying Port Security Settings.

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

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

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

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

```
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 14-2. 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 : Static           Address Limit : 1
  Action : Send Alarm

  Authorized Addresses
  -----
  00906d-fdcc00
```

Figure 14-3. Example of the Port Security Configuration Display for a Single Port

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

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

Listing Authorized and Detected MAC Addresses.

Syntax: show mac-address [port-list | mac-address | vlan < vid >]

Without an optional parameter, show mac-address lists the authorized MAC addresses that the switch detects on all ports.

mac-address: Lists the specified MAC address with the port on which it is detected as an authorized address.

port list: Lists the authorized MAC addresses detected on the specified port(s).

vlan < vid >: Lists the authorized MAC addresses detected on ports belonging to the specified VLAN.

```
ProCurve (config)# show mac-address
Status and Counters - Port Address Table
  MAC Address   Located on Port
  -----
0004ea-84d980 7
0004ea-84d9ee 7
000a57-4d8d40 5
      :
      :
00a0c9-f1786f 5

ProCurve (config)# show mac-address 7
Status and Counters - Port Address Table - 7
  MAC Address
  -----
0004ea-84d980
0004ea-84d9ee

ProCurve (config)# show mac-address 000a57-4d8d40
Status and Counters - Address Table - 000a57-4d8d40
  MAC Address : 000a57-4d8d40
  Located on Port : 5

ProCurve (config)# show mac-address vlan 1
Status and Counters - Address Table - VLAN 1
  MAC Address   Located on Port
  -----
0004ea-84d980 7
0004ea-84d9ee 7
000a57-4d8d40 5
      :
      :
00a0c9-f1786f 5
```

Figure 14-4. Examples of Show Mac-Address Outputs

Configuring Port Security

Using the CLI, you can:

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

Syntax: port-security

```
[e] <port-list>< learn-mode | address-limit | mac-address | action |  
clear-intrusion-flag >
```

< port-list >: Specifies a list of one or more ports to which the port-security command applies.

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

For the specified port:

- Identifies the method for acquiring authorized addresses.
- On switches covered in this guide, automatically invokes eavesdrop protection. (Refer to “Eavesdrop Protection” on page 14-5.)

continuous (Default): Appears in the factory-default setting or when you execute **no port-security**. Allows the port to learn addresses from the device(s) to which it is connected. In this state, the port accepts traffic from any device(s) to which it is connected. Addresses learned in the learn continuous mode will “age out” and be automatically deleted if they are not used regularly. The default age time is five minutes.

Addresses learned this way appear in the switch and port address tables and age out according to the **MAC Age Interval** in the System Information configuration screen of the Menu interface or the **show system-information** listing. You can set the MAC age out time using the CLI, SNMP, Web, or menu interfaces. For more information on the **mac-age-time** command refer to the chapter titled “Interface Access and System Information” in the Management and Configuration Guide for your switch.

— Continued —

Syntax: port-security (*Continued*)

learn-mode < continuous | static | port-access | configured | limited-continuous > (*Continued*)

static: *Enables you to use the **mac-address** parameter to specify the MAC addresses of the devices authorized for a port, and the **address-limit** parameter (explained below) to specify the number of MAC addresses authorized for the port. You can authorize specific devices for the port, while still allowing the port to accept other, non-specified devices until the device limit has been reached. That is, if you enter fewer MAC addresses than you authorized, the port authorizes the remaining addresses in the order in which it automatically learns them.*

*For example, if you use **address-limit** to specify three authorized devices, but use **mac-address** to specify only one authorized MAC address, the port adds the one specifically authorized MAC address to its authorized-devices list and the first two additional MAC addresses it detects.*

If, for example:

*You use **mac-address** to authorize MAC address 0060b0-880a80 for port A4.*

*You use **address-limit** to allow three devices on port A4 and the port detects these MAC addresses:*

- 1. 080090-1362f2*
- 2. 00f031-423fc1*
- 3. 080071-0c45a1*
- 4. 0060b0-880a80 (the address you authorized with the **mac-address** parameter)*

In this example port A4 would assume the following list of authorized addresses:

080090-1362f2 (the first address the port detected)

00f031-423fc1 (the second address the port detected)

*0060b0-880a80 (the address you authorized with the **mac-address** parameter)*

*The remaining MAC address detected by the port, 080071-0c45a1, is not allowed and is handled as an intruder. Learned addresses that become authorized do **not** age-out. See also “Retention of Static Addresses” on page 14-18.*

— Continued —

Syntax: port-security (*Continued*)

learn-mode < continuous | static | port-access | configured | limited-continuous > (*Continued*)

Caution: Using the **static** parameter with a device limit greater than the number of MAC addresses specified with **mac-address** can allow an un-wanted device to become “authorized”. This is because the port, to fulfill the number of devices allowed by the **address-limit** parameter (see below), automatically adds devices it detects until it reaches the specified limit.

Note: If 802.1X port-access is configured on a given port, then port-security learn-mode must be set to either **continuous** (the default) or **port-access**.

port-access: Enables you to use Port Security with (802.1X) Port-Based Access Control. Refer to chapter 13, Configuring Port-Based and User-Based Access Control (802.1X).

configured: Must specify which MAC addresses are allowed for this port. Range is 1 (default) to 8 and addresses are not ageable. Addresses are saved across reboots.

limited-continuous: Also known as MAC Secure, or “limited” mode. The limited parameter sets a finite limit to the number of learned addresses allowed per port. (You can set the range from 1, the default, to a maximum of 32 MAC addresses which may be learned by each port.)

All addresses are ageable, meaning they are automatically removed from the authorized address list for that port after a certain amount of time. Limited mode and the address limit are saved across reboots, but addresses which had been learned are lost during the reboot process. Addresses learned in the limited mode are normal addresses learned from the network until the limit is reached, but they are not configurable. (You cannot enter or remove these addresses manually if you are using **learn-mode** with the **limited-continuous** option.)

—Continued—

Syntax: port-security (*Continued*)

*Addresses learned this way appear in the switch and port address tables and age out according to the **MAC Age Interval** in the System Information configuration screen of the Menu interface or the **show system-information** listing. You can set the MAC age out time using the CLI, SNMP, Web, or menu interfaces. For more on the **mac-age-time** command, refer to the chapter titled “Interface Access and System Information” in the Management and Configuration Guide for your switch. To set the learn-mode to **limited** use this command syntax:*

```
port-security <port-list> learn-mode limited address-limit  
< 1..32 > action < none | send-alarm | send-disable >
```

*The default address-limit is **1** but may be set for each port to learn up to 32 addresses. The default action is **none**. To see the list of learned addresses for a port use the command:*

```
show mac < port-list >
```

```
address-limit < integer >
```

*When **learn-mode** is set to **static**, **configured**, or **limited-continuous**, the **address-limit** parameter specifies how many authorized devices (MAC addresses) to allow. Range: 1 (the default) to 8 for static and configured modes. For **learn-mode** with the **limited-continuous** option, the range is 1-32 addresses.*

—Continued—

Syntax: port-security (*Continued*)

mac-address [*<mac-addr>*] [*<mac-addr>*] . . . [*<mac-addr>*]

Available for **learn-mode** with the, **static**, **configured**, or **limited-continuous** option. Allows up to eight authorized devices (MAC addresses) per port, depending on the value specified in the **address-limit** parameter. The **mac-address limited** mode allows up to 32 authorized MAC addresses per port.

If you use **mac-address** with **static**, but enter fewer devices than you specified in the **address-limit** field, the port accepts not only your specified devices, but also as many other devices as it takes to reach the device limit. For example, if you specify four devices, but enter only two MAC addresses, the port will accept the first two non-specified devices it detects, along with the two specifically authorized devices. Learned addresses that become authorized do **not** age-out. See also “Retention of Static Addresses” on page 14-18.

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

Specifies whether an SNMP trap is sent to a network management station when Learn Mode is set to **static** and the port detects an unauthorized device, or when Learn Mode is set to continuous and there is an address change on a port.

none: Prevents an SNMP trap from being sent. **none** is the default value.

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

send-disable: Sends alarm and disables the port. Available only in the **static**, **port-access**, **configured**, or **limited learn** modes. Causes the switch to send an SNMP trap to a network management station and disable the port. If you subsequently re-enable the port without clearing the port’s intrusion flag, the port blocks further intruders, but the switch will not disable the port again until you reset the intrusion flag. See the Note on 14-36.

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

—Continued—

Syntax: port-security (*Continued*)

clear-intrusion-flag

Clears the intrusion flag for a specific port. (See “Reading Intrusion Alerts and Resetting Alert Flags” on page 14-34.)

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

Removes the specified learned MAC address(es) from the specified port.

Retention of Static Addresses

Static MAC addresses do not age-out. MAC addresses learned by using **learn-mode continuous** or **learn-mode limited-continuous** age out according to the currently configured MAC age time. (For information on the **mac-age-time** command, refer to the chapter titled “Interface Access and System Information” in the *Management and Configuration Guide* for your switch.

Learned Addresses. In the following two cases, a port in Static learn mode retains a learned MAC address even if you later reboot the switch or disable port security for that port:

- The port learns a MAC address after you configure the port for Static learn mode in both the startup-config file and the running-config file (by executing the **write memory** command).
- The port learns a MAC address after you configure the port for Static learn mode in only the running-config file and, after the address is learned, you execute **write memory** to configure the startup-config file to match the running-config file.

To remove an address learned using either of the preceding methods, do one of the following:

- Delete the address by using **no port-security <port-number> mac-address <mac-addr>**.
- Download a configuration file that does not include the unwanted MAC address assignment.
- Reset the switch to its factory-default configuration.

Assigned/Authorized Addresses. : If you manually assign a MAC address (using **port-security <port-number> address-list <mac-addr>**) and then execute **write memory**, the assigned MAC address remains in memory until you do one of the following:

- Delete it by using **no port-security <port-number> mac-address <mac-addr>**.
- Download a configuration file that does not include the unwanted MAC address assignment.
- Reset the switch to its factory-default configuration.

Specifying Authorized Devices and Intrusion Responses. This example configures port 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, but allow the intruder access to the network.

```
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 the switch is reset to its factory-default configuration. You can “turn off” authorized devices on a port by configuring the port to continuous Learn Mode, but subsequently reconfiguring the port to static Learn Mode restores those authorized devices.

Adding an Authorized Device to a Port. To simply add a device (MAC address) to a port's existing Authorized Addresses list, enter the port number with the **mac-address** parameter and the device's MAC address. *This assumes that Learn Mode is set to **static** and the Authorized Addresses list is not full* (as determined by the current Address Limit value). For example, suppose port A1 allows two authorized devices, but has only one device in its Authorized Address list:

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

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

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

The Address Limit has not been reached.

Figure 14-5. 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 would be:

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

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

The Address Limit has been reached.

Figure 14-6. Example of Adding a Second Authorized Device to a Port

(The message **Inconsistent value** appears if the new MAC address exceeds the current Address Limit or specifies a device that is already on the list. Note that if you change a port from static to continuous learn mode, the port retains in memory any authorized addresses it had while in static mode. If you 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:

```
HPswitch(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 14-7. 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. This command option removes unwanted devices (MAC addresses) from the Authorized Addresses list. (An Authorized Address list is available for each port for which Learn Mode is currently set to “Static”. Refer to the command syntax listing under “Configuring Port Security” on page 14-12.)

Caution

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

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

Note

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

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

```
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
```

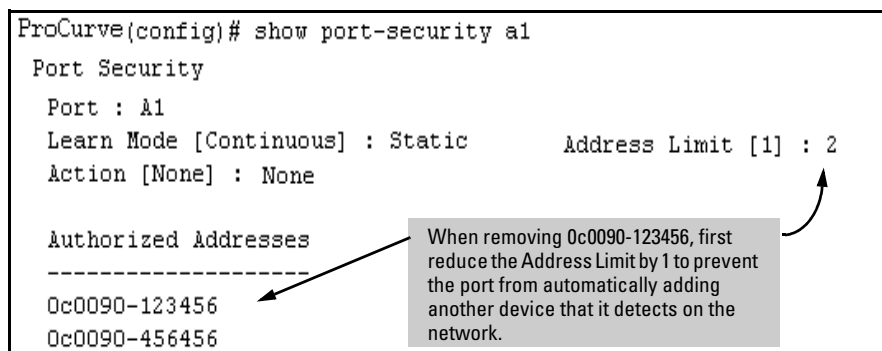


Figure 14-8. Example of Two Authorized Addresses on Port A1

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 14-9. 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 “meshing” or Spanning Tree Protocol (STP) to speed up network performance by providing multiple paths for devices, using MAC Lockdown either will not work or else it defeats the purpose of having multiple data paths.

The purpose of using MAC Lockdown is to prevent a malicious user from “hijacking” an approved MAC address so they can steal data traffic being sent to that address.

As we have seen, MAC Lockdown can help prevent this type of hijacking by making sure that all traffic to a specific MAC address goes only to the proper port on a switch which is supposed to be connected to the real device bearing that MAC address.

However, you can run into trouble if you incorrectly try to deploy MAC Lockdown in a network that uses multiple path technology, like Spanning Tree or “mesh networks.”

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

Configuring and Monitoring Port Security MAC Lockdown

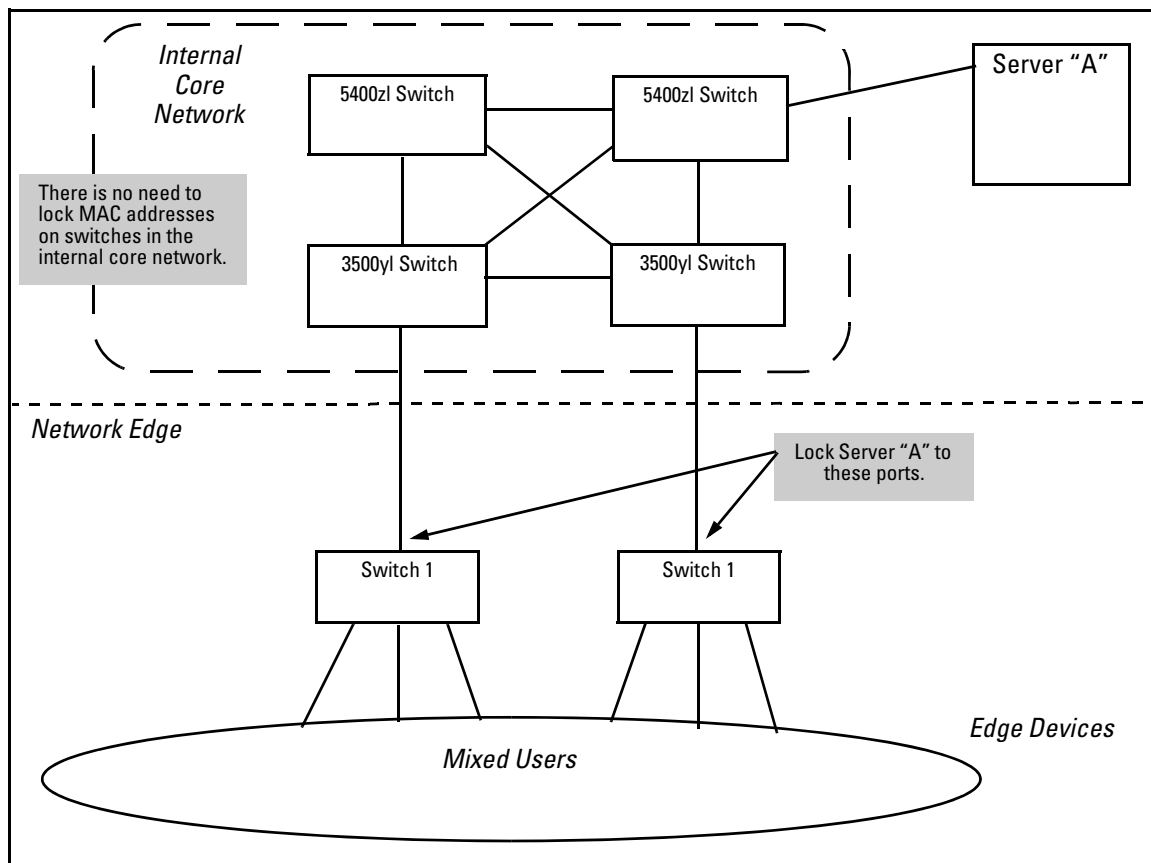


Figure 14-10. MAC Lockdown Deployed At the Network Edge Provides Security

Basic MAC Lockdown Deployment. In the Model Network Topology shown above, the switches that are connected to the edge of the network each have one and only one connection to the core network. This means each switch has only one path by which data can travel to Server A. You can use MAC Lockdown to specify that all traffic intended for Server A's MAC Address must go through the one port on the edge switches. That way, users on the edge can still use other network resources, but they cannot "spoof" Server A and hijack data traffic which is intended for that server alone.

The key points for this Model Topology are:

- The Core Network is separated from the edge by the use of switches which have been “locked down” for security.
- All switches connected to the edge (outside users) each have only one port they can use to connect to the Core Network and then to Server A.
- Each switch has been configured with MAC Lockdown so that the MAC Address for Server A has been locked down to one port per switch that can connect to the Core and Server A.

Using this setup Server A can be moved around within the core network, and yet MAC Lockdown will still prevent a user at the edge from hijacking its address and stealing data.

Please note that in this scenario a user with bad intentions at the edge can still “spoof” the address for Server A and send out data packets that look as though they came from Server A. The good news is that because MAC Lockdown has been used on the switches on the edge, any traffic that is sent *back* to Server A will be sent to the proper MAC Address because MAC Lockdown has been used. The switches at the edge will not send Server A’s data packets anywhere but the port connected to Server A. (Data would not be allowed to go beyond the edge switches.)

Caution

Using MAC Lockdown still does not protect against a hijacker *within the core!* In order to protect against someone spoofing the MAC Address for Server A inside the Core Network, you would have to lock down each and every switch inside the Core Network as well, not just on the edge.

Problems Using MAC Lockdown in Networks With Multiple Paths. Now let’s take a look at a network topology in which the use of MAC Lockdown presents a problem. In the next figure, Switch 1 (on the bottom-left) is located at the edge of the network where there is a mixed audience that might contain hackers or other malicious users. Switch 1 has two paths it could use to connect to Server A. If you try to use MAC Lockdown here to make sure that all data to Server A is “locked down” to one path, connectivity problems would be the result since both paths need to be usable in case one of them fails.

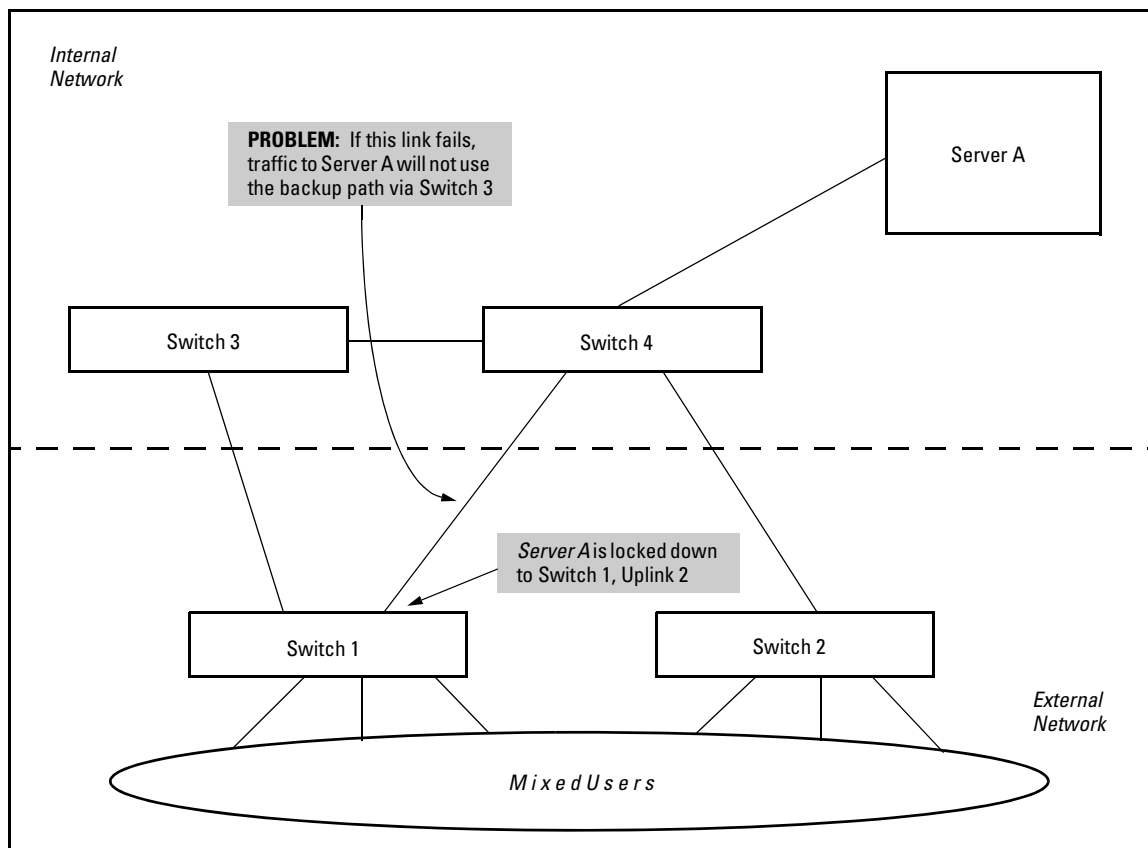


Figure 14-11. Connectivity Problems Using MAC Lockdown with Multiple Paths

The resultant connectivity issues would prevent you from locking down Server A to Switch 1. And when you remove the MAC Lockdown from Switch 1 (to prevent broadcast storms or other connectivity issues), you then open the network to security problems. The use of MAC Lockdown as shown in the above figure would defeat the purpose of using MSTP or having an alternate path.

Technologies such as MSTP or “meshing” are primarily intended for an internal campus network environment in which all users are trusted. MSTP and “meshing” do not work well with MAC Lockdown.

If you deploy MAC Lockdown as shown in the Model Topology in figure 14-10 (page 14-28), you should have no problems with either security or connectivity.

MAC Lockout

MAC Lockout involves configuring a MAC address on all ports and VLANs for a switch so that any traffic to or from the “locked-out” MAC address will be dropped. This means that all data packets addressed to or from the given address are stopped by the switch. MAC Lockout is implemented on a per switch assignment.

You can think of MAC Lockout as a simple blacklist. The MAC address is locked out on the switch and on all VLANs. No data goes out or in from the blacklisted MAC address to a switch using MAC Lockout.

To fully lock out a MAC address from the network it would be necessary to use the MAC Lockout command on all switches.

To use MAC Lockout you must first know the MAC Address you wish to block.

Syntax: [no] lockout-mac < mac-address >

How It Works. Let’s say a customer knows there are unauthorized wireless clients who should not have access to the network. The network administrator “locks out” the MAC addresses for the wireless clients by using the MAC Lockout command (**lockout-mac <mac-address>**). When the wireless clients then attempt to use the network, the switch recognizes the intruding MAC addresses and prevents them from sending or receiving data on that network.

If a particular MAC address can be identified as unwanted on the switch then that MAC Address can be disallowed on all ports on that switch with a single command. You don’t have to configure every single port—just perform the command on the switch and it is effective for all ports.

MAC Lockout overrides MAC Lockdown, port security, and 802.1X authentication.

You cannot use MAC Lockout to lock:

- Broadcast or Multicast Addresses (Switches do not learn these)
- Switch Agents (The switch's own MAC Address)

There are limits for the number of VLANs, Multicast Filters, and Lockout MACs that can be configured concurrently as all use MAC table entries. The limits are shown below.

Table 14-12. Limits on Lockout MACs

# VLANs	# Multicast Filters	# Lockout MACs
<= 1024	16	16
1025-2048	8	8

If someone using a locked out MAC address tries to send data through the switch a message is generated in the log file:

Lockout logging format:

```
W 10/30/03 21:35:15 maclock: 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 14-26.)

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 network management applications such as ProCurve Manager 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.

```
Status and Counters - Intrusion Log
Port  MAC Address          Date / Time
-----
A1    080009-e93d4f            03/07/06 21:09:34
A1    080009-e93d4f            03/07/06 10:18:43
```

Figure 14-13. Example of Multiple Intrusion Log Entries for the Same Port

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

Keeping the Intrusion Log Current by Resetting Alert Flags

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

Note on Send-Disable Operation

On a given port, if the intrusion action is to send an SNMP trap and then disable the port (**send-disable**), and an intruder is detected on the port, then the switch sends an SNMP trap, sets the port's alert flag, and disables the port. If you re-enable the port without resetting the port's alert flag, then the port operates as follows:

- The port comes up and will block traffic from unauthorized devices it detects.
- If the port detects another intruder, it will send another SNMP trap, but will not become disabled again unless you first reset the port's intrusion flag.

This operation enables the port to continue passing traffic for authorized devices while you take the time to locate and eliminate the intruder. Otherwise, the presence of an intruder could cause the switch to repeatedly disable the port.

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

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

- From the Main Menu select:

- 1. Status and Counters**
- 4. Port Status**

```

===== CONSOLE - MANAGER MODE =====
                        Status and Counters - Port Status
-----
Port      Type      Intrusion      Enabled  Status  Mode      Flow
Alert
-----
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.
    
```

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

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

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

```

===== CONSOLE - MANAGER MODE =====
                        Status and Counters - Intrusion Log
-----
Port      MAC Address      Date / Time      System Time of Intrusion on Port
-----
A3      080009-6563e2    08/08/02 16:58:02
A1      0060b0-896e00    08/08/02 15:28:21
A3      080009-cf558f    prior to 08/08/02 10:28:58

Actions->  Back      Reset alert flags  Help

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

MAC Address of Intruding Device on

System Time of Intrusion on Port

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

Figure 14-15. Example of the Intrusion Log Display

The example in Figure 7-11 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 14-14 on page 14-37) 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 alert flag for 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 14-15, 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 14-42.)

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

Intrusion Alert on port

Figure 14-16. Example of an Unacknowledged Intrusion Alert in a Port Status Display

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

```

ProCurve# show port-security intrusion-log
Status and Counters - Intrusion Log

```

Port	MAC Address	Date / Time
A1	080009-e93d4f	07/03/06 21:09:34
A1	080009-21ae84	07/03/06 17:26:27
A1	080009-e93d4f	prior to 07/03/06 17:18:43
	0 secs	
	0 secs	

MAC Address of latest Intruder on Port A1

Dates and Times of Intrusions

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

Figure 14-17. 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.

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

To clear the intrusion from port A1 and enable the switch to enter any subsequent intrusion for port A1 in the Intrusion Log, execute the port-security **clear-intrusion-flag** command. If you then re-display the port status screen, you will see that the Intrusion Alert entry for port A1 has changed to **“No”**. (Executing **show port-security intrusion-log** again will result in the same display as above, and does not include the Intrusion Alert status.)

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

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

Intrusion Alert on port A1 is now

Figure 14-18. Example of Port Status Screen After Alert Flags Reset

For more on clearing intrusions, see “Note on Send-Disable Operation” on page 14-36

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  ----
```

Figure 14-19. Example of Log Listing With and Without Detected Security Violations

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

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

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

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

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

Operating Notes for Port Security

Identifying the IP Address of an Intruder. The Intrusion Log lists detected intruders by MAC address. If you are using ProCurve Manager to manage your network, you can use the device properties page to link MAC addresses to their corresponding IP addresses.

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

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

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

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

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

LACP Not Available on Ports Configured for Port Security. To maintain security, LACP is not allowed on ports configured for port security. If you configure port security on a port on which LACP (active or passive) is configured, the switch removes the LACP configuration, displays a notice that LACP is disabled on the port(s), and enables port security on that port. For example:

```
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.

— This page is intentionally unused —

Using Authorized IP Managers

Contents

Overview	15-2
Options	15-3
Access Levels	15-3
Defining Authorized Management Stations	15-4
Overview of IP Mask Operation	15-4
Menu: Viewing and Configuring IP Authorized Managers	15-5
CLI: Viewing and Configuring Authorized IP Managers	15-6
Listing the Switch's Current Authorized IP Manager(s)	15-6
Configuring IP Authorized Managers for the Switch	15-7
Web: Configuring IP Authorized Managers	15-9
Building IP Masks	15-9
Configuring One Station Per Authorized Manager IP Entry	15-9
Configuring Multiple Stations Per Authorized Manager IP Entry ..	15-10
Additional Examples for Authorizing Multiple Stations	15-12
Operating Notes	15-12

Overview

Authorized IP Manager Features

Feature	Default	Menu	CLI	Web
Listing (Showing) Authorized Managers	n/a	page 15-5	page 15-6	page 15-9
Configuring Authorized IP Managers	None	page 15-5	page 15-6	page 15-9
Building IP Masks	n/a	page 15-9	page 15-9	page 15-9
Operating and Troubleshooting Notes	n/a	page 15-12	page 15-12	page 15-12

The Authorized IP Managers feature uses IP addresses and masks to determine which stations (PCs or workstations) can access the switch through the network. This covers access through the following means:

- Telnet and other terminal emulation applications
- The switch’s web browser interface
- SNMP (with a correct community name)

Also, when configured in the switch, the Authorized IP Managers feature takes precedence over local passwords, TACACS+, RADIUS, Port-Based Access Control (802.1X), and Port Security. This means that the IP address of a networked management device must be authorized before the switch will attempt to authenticate the device by invoking any other access security features. If the Authorized IP Managers feature disallows access to the device, then access is denied. Thus, with authorized IP managers configured, having the correct passwords is not sufficient for accessing the switch through the network unless the station attempting access is also included in the switch’s Authorized IP Managers configuration.

You can use Authorized IP Managers along with other access security features to provide a more comprehensive security fabric than if you use only one or two security options.

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 unauthorized 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 15-9.)
- **Authorizing Multiple Stations:** The table entry uses the IP Mask to authorize access to the switch from a defined group of stations. This is useful if you want to easily authorize several stations to have access to the switch without having to type in an entry for every station. All stations in the group defined by the one Authorized Manager IP table entry and its associated IP mask will have the same access level—Manager or Operator. (For more on this topic, refer to “Configuring Multiple Stations Per Authorized Manager IP Entry” on page 15-10.)

To configure the switch for authorized manager access, enter the appropriate *Authorized Manager IP* value, specify an *IP Mask*, and select either **Manager** or **Operator** for the *Access Level*. The IP Mask determines how the Authorized Manager IP value is used to allow or deny access to the switch by a management station.

Overview of IP Mask Operation

The default IP Mask is 255.255.255.255 and allows switch access only to a station having an IP address that is identical to the Authorized Manager IP parameter value. (“255” in an octet of the mask means that only the exact value in the corresponding octet of the Authorized Manager IP parameter is allowed in the IP address of an authorized management station.) However, you can alter the mask and the Authorized Manager IP parameter to specify ranges of authorized IP addresses. For example, a mask of **255.255.255.0** and any value for the Authorized Manager IP parameter allows a range of 0 through 255 in the 4th octet of the authorized IP address, which enables a block of up to 254 IP addresses for IP management access (excluding 0 for the network and 255 for broadcasts). A mask of **255.255.255.252** uses the 4th octet of a given Autho-

rized Manager IP address to authorize four IP addresses for management station access. The details on how to use IP masks are provided under “Building IP Masks” on page 15-9.

Note

The IP Mask is a method for recognizing whether a given IP address is authorized for management access to the switch. This mask serves a different purpose than IP subnet masks and is applied in a different manner.

Menu: Viewing and Configuring IP Authorized Managers

From the console Main Menu, select:

2. Switch Configuration ...

7. IP Authorized Managers

```
----- TELNET - MANAGER MODE -----  
Switch Configuration - IP Managers  
  
Authorized Manager IP      IP Mask      Access Level  
-----  
13.28.227.101             255.255.255.252  Manager  
13.28.227.104             255.255.255.254  Manager  
13.28.227.106             255.255.255.0    Operator  
13.28.227.125             255.255.255.255  Manager  
  
Actions->  Back  Add  Edit  Delete  Help  
  
Return to previous screen.  
Use up/down arrow keys to change record selection, left/right arrow keys to  
change action selection, and <Enter> to execute action.
```

Figure 15-1. Example of How To Add an Authorized Manager Entry

Using Authorized IP Managers

Defining Authorized Management Stations

```

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

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

Actions->  _Cancel    _Edit    _Save    _Help

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

```

1. Enter an Authorized Manager IP address here.

2. Use the default mask to allow access by one management device, or edit the mask to allow access by a block of management devices. See "Building IP Masks" on page 15-9.

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

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

5. Applies only to access through Telnet, SNMPv1, and SNMPv2c. Refer to the note on page 15-3.

Figure 15-2. 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 15-1), 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	15-7
<ip-address>	15-8
mask <mask-bits>	15-8
<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 15-3. 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 15-3.

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

```

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
  
```

Using Authorized IP Managers

Defining Authorized Management Stations

If you omit the *<mask bits>* when adding a new authorized manager, the switch automatically uses **255.255.255.255**. If you do not specify either Manager or Operator access, the switch assigns the Manager access. For example:

```
ProCurve (config)# ip authorized-managers [10.28.227.105]
ProCurve (config)# show ip authorized-managers
```

IP Managers

Authorized Manager IP	IP Mask	Access Level
10.28.227.105	255.255.255.255	Manager

Omitting a mask in the ip authorized-managers command results in a default mask of 255.255.255.255, which authorizes only the specified station. Refer to "Configuring Multiple Stations Per Authorized Manager IP Entry" on page 15-10.

Figure 15-4. Example of Specifying an IP Authorized Manager with the Default Mask

To Edit an Existing Manager Access Entry. To change the mask or access level for an existing entry, use the entry's IP address and enter the new value(s). (Notice that any parameters not included in the command will be set to their default.):

```
ProCurve (config)# ip authorized-managers
    10.28.227.101 255.255.255.0 access operator
```

The above command replaces the existing mask and access level for IP address 10.28.227.101 with 255.255.255.0 and operator.

The following command replaces the existing mask and access level for IP address 10.28.227.101 with 255.255.255.255 and manager (the defaults) because the command does not specify either of these parameters.

```
ProCurve (config)# ip authorized-managers 10.28.227.101
```

To Delete an Authorized Manager Entry. This command uses the IP address of the authorized manager you want to delete:

```
ProCurve (config)# no ip authorized-managers 10.28.227.101
```

Web: Configuring IP Authorized Managers

In the web browser interface you can configure IP Authorized Managers as described below.

To Add, Modify, or Delete an IP Authorized Manager address:

1. Click on the **Security** tab.
2. Click on [Authorized Addresses].
3. Enter the appropriate parameter settings for the operation you want.
4. Click on [Add], [Replace], or [Delete] to implement the configuration change.

For web-based help on how to use the web browser interface screen, click on the [?] button provided on the web browser screen.

Building IP Masks

The IP Mask parameter controls how the switch uses an Authorized Manager IP value to recognize the IP addresses of authorized manager stations on your network.

Configuring One Station Per Authorized Manager IP Entry

This is the easiest way to apply a mask. If you have ten or fewer management and/or operator stations, you can configure them by adding the address of each to the Authorized Manager IP list with **255.255.255.255** for the corresponding mask. For example, as shown in figure 15-3 on page 15-7, if you configure an IP address of **10.28.227.125** with an IP mask of **255.255.255.255**, only a station having an IP address of **10.28.227.125** has management access to the switch.

Figure 15-5. 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.

Figure 15-6. Analysis of IP Mask for Multiple-Station Entries

	1st Octet	2nd Octet	3rd Octet	4th Octet	Manager-Level or Operator-Level Device Access
IP Mask	255	255	255	0	The "255" in the first three octets of the mask specify that only the exact value in the octet of the corresponding IP address is allowed. However, the zero (0) in the 4th octet of the mask allows any value between 0 and 255 in that octet of the corresponding IP address. This mask allows switch access to any device having an IP address of 10.28.227.xxx, where xxx is any value from 0 to 255.
Authorized Manager IP	10	28	227	125	
IP Mask	255	255	255	249	In this example (figure 15-7, 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	

Figure 15-7. Example of How the Bitmap in the IP Mask Defines Authorized Manager Addresses

4th Octet of IP Mask:		249							
4th Octet of Authorized IP Address:		5							
Bit Numbers	Bit 7	Bit 6	Bit 5	Bit 4	Bit 3	Bit 2	Bit 1	Bit 0	
Bit Values	128	64	32	16	8	4	2	1	
4th Octet of IP Mask (249)									Bits 1 and 2 in the mask are "off", and bits 0 and 3 - 7 are "on", creating a value of 249 in the 4th octet. Where a mask bit is "on", the corresponding bit setting in the address of a potentially authorized station must match the IP Authorized Address setting for that same bit. Where a mask bit is "off" the corresponding bit setting in the address can be either "on" or "off". In this example, in order for a station to be authorized to access the switch: <ul style="list-style-type: none"> • The first three octets of the station's IP address must match the Authorized IP Address. • Bit 0 and Bits 3 through 6 of the 4th octet in the station's address must be "on" (value = 1). • Bit 7 of the 4th octet in the station's address must be "off" (value = 0). • Bits 1 and 2 can be either "on" or "off". This means that stations with the IP address 13.28.227.X (where X is 121, 123, 125, or 127) are authorized.
4th Octet of IP Authorized Address (125)									

Additional Examples for Authorizing Multiple Stations

	Entries for Authorized Manager List	Results
IP Mask	255 255 0 255	This combination specifies an authorized IP address of 10.33.xxx.1. It could be applied, for example, to a subnetted network where each subnet is defined by the third octet and includes a management station defined by the value of "1" in the fourth octet of the station's IP address.
Authorized Manager IP	10 33 248 1	
IP Mask	255 238 255 250	Allows 230, 231, 246, and 247 in the 2nd octet, and 194, 195, 198, 199 in the 4th octet.
Authorized Manager IP	10 247 100 195	

Operating Notes

- **Network Security Precautions:** You can enhance your network's security by keeping physical access to the switch restricted to authorized personnel, using the password features built into the switch, using the additional security features described in this manual, and preventing unauthorized access to data on your management stations.
- **Modem and Direct Console Access:** Configuring authorized IP managers does not protect against access to the switch through a modem or direct Console (RS-232) port connection.
- **Duplicate IP Addresses:** If the IP address configured in an authorized management station is also configured (or "spoofed") in another station, the other station can gain management access to the switch even though a duplicate IP address condition exists.
- **Web Proxy Servers:** If you use the web browser interface to access the switch from an authorized IP manager station, it is recommended that you avoid the use of a web proxy server in the path between the station and the switch. This is because switch access through a web proxy server requires that you first add the web proxy server to the Authorized Manager IP list. *This reduces security by opening switch access to anyone who uses the web proxy server.* The following two options outline how to eliminate a web proxy server from the path between a station and the switch:

- Even if you need proxy server access enabled in order to use other applications, you can still eliminate proxy service for web access to the switch. To do so, add the IP address or DNS name of the switch to the non-proxy, or “Exceptions” list in the web browser interface you are using on the authorized station.
- If you don’t need proxy server access at all on the authorized station, then just disable the proxy server feature in the station’s web browser interface.

— This page is intentionally unused —

Key Management System

Contents

Overview	16-2
Terminology	16-2
Configuring Key Chain Management	16-3
Creating and Deleting Key Chain Entries	16-3
Assigning a Time-Independent Key to a Chain	16-4
Assigning Time-Dependent Keys to a Chain	16-5

Overview

The switches covered in this guide provide support for advanced routing capabilities. Security turns out to be extremely important as complex networks and the internet grow and become a part of our daily life and business. This fact forces protocol developers to improve security mechanisms employed by their protocols, which in turn becomes an extra burden for system administrators who have to set up and maintain them. One possible solution to the problem is to centralize the mechanisms used to configure and maintain security information for all routing protocols. The Key Management System (KMS) can carry this burden.

KMS is designed to configure and maintain key chains. A key chain is a set of keys with a timing mechanism for activating and deactivating individual keys. KMS provides specific instances of routing protocols with one or more Send or Accept keys that must be active at the time of a request. A *protocol instance* is usually an interface on which the protocol is running.

Feature	Default	Menu	CLI	Web
Generating a Key Chain	n/a	n/a	page 16-3	n/a
Generating a Time-Independent key	n/a	n/a	page 16-4	n/a
Generating a Time-Dependent key	n/a	n/a	page 16-5	n/a

Terminology

- **Key Chain:** A key or set of keys assigned for use by KMS-enabled protocols. A key chain may optionally contain the time to activate and deactivate a particular key.
- **Time-Independent Key:** A key that has no activate or deactivate time associated with it. This type of key does not expire, which eliminates the need for a key chain.
- **Time-Dependent key:** a key that has an activate and deactivate time associated with the Accept and Send processes. Time-Dependent keys expire, which means a key chain is needed to keep the assigned protocols supplied with keys.
- **Key Management System (KMS) Enabled Protocol:** A protocol that uses KMS to store authentication key information.

Configuring Key Chain Management

KMS-Related CLI Commands in This Section	Page
show key-chain < chain_name >	page 16-3
[no] key-chain chain_name	page 16-3
[no] key-chain chain_name key Key_ID	page 16-4

The Key Management System (KMS) has three configuration steps:

1. Create a key chain entry.
2. Assign a time-independent key or set of time-dependent keys to the Key Chain entry. The choice of key type is based on the level of security required for the protocol to which the key entry will be assigned.
3. Assign the key chain to a KMS-enabled protocol.

This procedure is protocol-dependent. For information on a specific protocol, refer to the chapter covering that protocol in the *Management and Configuration Guide* for your switch.

Creating and Deleting Key Chain Entries

To use the Key Management System (KMS), you must create one or more key chain entries. An entry can be the pointer to a single time-independent key or a chain of time-dependent keys

Syntax: [no] key-chain < chain_name >

*Generate or delete a key chain entry. Using the optional **no** form of the command deletes the key chain. The < chain_name > parameter can include up to 32 characters.*

show key-chain

Displays the current key chains on the switch and their overall status.

For example, to generate a new key chain entry:

```
ProCurve.(config)# key-chain Procurvel
ProCurve.(config)# show key-chain
```

Chain Name	Keys	Active	Expired
Procurvel	0	0	0

Annotations:
- Arrow from 'key-chain Procurvel' points to 'Add new key chain Entry "Procurve1".'
- Arrow from 'show key-chain' points to 'Display key chain entries.'

Figure 16-1. Adding a New Key Chain Entry

After you add an entry, you can assign key(s) to it for use by a KMS-enabled protocol.

Assigning a Time-Independent Key to a Chain

A time-independent key has no Accept or Send time constraints. It is valid from boot-up until you change it. If you use a time-independent key, then it is the only key needed for a key chain entry.

Syntax: [no] key-chain < chain_name > key < key_id >

*Generates or deletes a key in the key chain entry < chain_name >. Using the optional **no** form of the command deletes the key. The < key_id > is any number from 0-255.*

[key-string < key_str >]

This option lets you specify the key value for the protocol using the key. The < key_str > can be any string of up to 14 characters in length.

[accept-lifetime infinite] [send-lifetime infinite]

accept-lifetime infinite: *Allows packets with this key to be accepted at any time from boot-up until the key is removed.*

send-lifetime infinite: *Allows the switch to send this key as authorization, from boot-up until the key is removed.*

show key-chain < chain_name >

Displays the detail information about the keys used in the key chain named < chain_name >.

For example, to generate a new time-independent key for the Procurvel key chain entry:


```

ProCurve (config)# key-chain Procurvel key 1
ProCurve (config)# show key-chain Procurvel

```

Chain - Procurvel

Key	Accept Start GMT	Accept Stop GMT	Send Start GMT	Send Stop GMT
1	Bootup	Infinite	Bootup	Infinite

OSPF Interface References

Interface

OSPF Virtual Link References

← Adds a new Time-Independent key to the "Procurve1" chain.

← Displays keys in the key chain entry.

Figure 16-2. Example of Adding and Displaying a Time-Independent Key to a Key Chain Entry

Assigning Time-Dependent Keys to a Chain

A time-dependent key has Accept or Send time constraints. It is valid only during the times that are defined for the key . If a time-dependent key is used, there is usually more than one key in the key chain entry.

Syntax: [no] key-chain < chain_name > key < key_id >

*Generates or deletes a key in the key chain entry < chain_name >. Using the optional **no** form of the command deletes the key. The < key_id > is any number from 0-255.*

[key-string < key_str >]

This option specifies the key value referenced by the protocol using the key. The < key_str > can be any string up to 14 characters in length.

accept-lifetime < mm/dd/yy [yy] hh:mm:ss | now >

*Specifies the **start** date and time of the valid period in which the switch can use this key to authenticate inbound packets.*

duration < mm/dd/yy[yy] hh:mm:ss | seconds >

*Specifies the **time period** during which the switch can use this key to authenticate inbound packets. Duration is either an end date and time or the number of seconds to allow after the start date and time (which is the **accept-lifetime** setting).*

send-lifetime <mm/dd/yy[yy] hh:mm:ss | now>

*Specifies the **start** date and time of the valid period in which the switch can transmit this key as authentication for outbound packets.*

duration < mm/dd/yy[yy] hh:mm:ss | seconds >

*Specifies the **time period** during which the switch can use this key to authenticate outbound packets. Duration is either an end date and time or the number of seconds to allow after the start date and time (which is the **accept-lifetime** setting).*

show key-chain < chain_name >

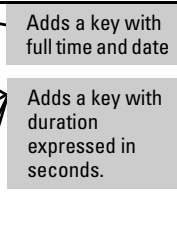
Displays the detail information about the keys used in the key chain named < chain_name >.

Note

Using time-dependent keys requires that all the switches have accurate, synchronized time settings. You can manually set the time or use the Time protocol feature included in the switches. For more information, refer to the chapter covering time protocols in the *Management and Configuration Guide* for your switch.

For example, to add a number of keys to the key chain entry “Procurve2”:

```
ProCurve (config)# key-chain Procurve2 key 1 accept-lifetime 01/17/03 8:00:00
01/18/03 8:10:00 send-lifetime 01/17/03 8:00:00 01/18/03 8:00:00
ProCurve (config)# key-chain Procurve2 key 2 accept-lifetime 01/18/03 8:00:00
duration 87000 send-lifetime 01/18/03 8:00:00 duration 86400
ProCurve (config)# key-chain Procurve2 key 3 accept-lifetime 01/19/03 8:00:00
duration 87000 send-lifetime 01/19/03 8:00:00 duration 86400
ProCurve (config)# key-chain Procurve2 key 4 accept-lifetime 01/20/03 8:00:00
duration 87000 send-lifetime 01/20/03 8:00:00 duration 86400
ProCurve (config)# key-chain Procurve2 key 5 accept-lifetime 01/21/03 8:00:00
duration 87000 send-lifetime 01/21/03 8:00:00 duration 86400
```



← Adds a key with full time and date

← Adds a key with duration expressed in seconds.

Figure 16-3. Adding Time-Dependent Keys to a Key Chain Entry

Note

Given transmission delays and the variations in the time value from switch to switch, it is advisable to include some flexibility in the Accept lifetime of the keys you configure. Otherwise, the switch may disregard some packets because either their key has expired while in transport or there are significant time variations between switches.

To list the result of the commands in figure 16-3:

```
ProCurve(config)# show key-chain Procurve2

Chain - Procurve2

Key | Accept Start GMT   Accept Stop GMT    Send Start GMT     Send Stop GMT
---+-----
 1  | 01/17/03 08:00:00 01/18/03 08:10:00 01/17/03 08:00:00 01/18/03 08:00:00
 2  | 01/18/03 08:00:00 01/19/03 08:10:00 01/18/03 08:00:00 01/19/03 08:00:00
 3  | 01/19/03 08:00:00 01/20/03 08:10:00 01/19/03 08:00:00 01/20/03 08:00:00
 4  | 01/20/03 08:00:00 01/21/03 08:10:00 01/20/03 08:00:00 01/21/03 08:00:00
 5  | 01/21/03 08:00:00 01/22/03 08:10:00 01/21/03 08:00:00 01/22/03 08:00:00

OSPF Interface References

Interface
-----

OSPF Virtual Link References

Area/Virtual Link
-----
```

Figure 16-4. Display of Time-Dependent Keys in the Key Chain Entry

You can use **show key-chain** to display the key status at the time the command is issued. Using the information from the example configuration in figures 16-3 and 16-4, if you execute **show key-chain** at 8:05 on 01/19/03, the display would appear as follows:

```
ProCurve(config)# show key-chain

Key Chains

Chain Name                               Keys Active Expired
-----
Procurve1                                 1    1    0
Procurve2                                 5    2    1
```

Figure 16-5. Status of Keys in Key Chain Entry "Procurve2"

The “Procurve1” key chain entry is a time-independent key and will not expire. “Procurve2” uses time-dependent keys, which result in this data:

Expired = 1	Key 1 has expired because its lifetime ended at 8:10 on 01/18/03, the previous day.
Active = 2	Key 2 and 3 are both active for 10 minutes from 8:00 to 8:10 on 1/19/03.

Keys 4 and 5 are either not yet active or expired. The total number of keys is 5.

Index

Numerics

3DES ... 8-3, 9-3

802.1X

ACL, effect on ... 10-20

802.1X access control

authenticate users ... 13-5

authentication methods ... 13-4

authentication, local ... 13-6

authentication, user-based ... 13-4

authenticator ... 13-17

backend state ... 13-51

operation ... 13-9

show commands ... 13-48

state ... 13-51

unblock port ... 13-5

authorized-client VLAN, defined ... 13-6

auth-vid ... 13-21

auto ... 13-19

blocked port, trunked ... 13-13

caution, unauthorized-client VLAN ... 13-34

CHAP ... 13-3

chap-radius ... 13-21

clear-statistics ... 13-23

client, effect of disconnect ... 13-34

client-limit, no ... 13-18

configure

commands ... 13-16

displaying configuration ... 13-48

overview ... 13-15

port ... 13-17

configuring method ... 13-21

control all clients ... 13-12

control command ... 13-19

convert to port-based ... 13-18

CoS override ... 13-48, 13-51

counters ... 13-48

delay move to unauthorized-client

VLAN ... 13-30

delay Unauth-Client VLAN ... 13-21

DHCP server ... 13-35

displaying 802.1X port configuration ... 13-50

EAP ... 13-3

EAPOL ... 13-7, 13-49

eap-radius ... 13-21

enabling controlled directions ... 13-24

enabling on ports ... 13-17

enabling on switch ... 13-23

features ... 13-3

force authorized ... 13-19, 13-52

force unauthorized ... 13-19, 13-52

general setup ... 13-14

guest VLAN ... 13-7, 13-8, 13-27, 13-33

GVRP ... 13-56

effect ... 13-59

initialize ... 13-23

LACP not allowed ... 13-61

local ... 13-21

local username and password ... 13-4

logoff-period ... 13-21

max-requests ... 13-20

MD-5 ... 13-45

MD5 ... 13-7

meshing, not supported ... 13-13

messages802.1X access control

event log messages ... 13-61

multiple clients ... 13-34

multiple clients, same VLAN ... 13-4

open port ... 13-4

open VLAN

authorized client ... 13-28

configuration ... 13-38, 13-40

general operation ... 13-26

mode ... 13-26, 13-32

operating notes ... 13-41

operating rules ... 13-33

PVID, no ... 13-51

security breach ... 13-41

set up ... 13-37

status ... 13-50, 13-52

status, viewing ... 13-51

suspended VLAN ... 13-52

unauthorized client ... 13-28

use model ... 13-28

VLAN, after authentication ... 13-28, 13-34,
13-41

VLAN, tagged ... 13-28, 13-29, 13-34, 13-41,
13-54

overview ... 1-8, 13-3

port, supplicant ... 13-14

- port-based
 - access ... 13-4
 - client without authentication ... 13-5
 - effect of Web/MAC Auth client ... 13-60
 - enable ... 13-17, 13-43
 - latest client, effect ... 13-5
 - multiple client access ... 13-5
 - multiple clients authenticating ... 13-5
 - no client limit ... 13-4
 - open port ... 13-4
 - operation ... 13-5
 - recommended use ... 13-5
 - return to ... 13-18
 - See also* user-based.
 - single client authenticates ... 13-5
 - tagged VLAN membership ... 13-5
 - unauthorized client risk ... 13-5
 - untagged VLAN membership ... 13-5, 13-27
 - with Web/MAC authentication ... 13-6
- port-security ... 13-39
- port-security use ... 13-5
- port-security, with 802.1X ... 13-42
- priority of VLAN, per-port ... 13-10, 13-27
- PVID ... 13-53
- quiet-period ... 13-20
- RADIUS
 - effect on VLAN operation ... 13-56
- RADIUS ... 13-3
 - VLAN assignment ... 13-33
- RADIUS host IP address ... 13-22
- Rate-Limit override ... 13-51
- reauthenticate ... 13-23
- reauth-period ... 13-21
- rules of operation ... 13-12
- server-timeout ... 13-20
- show commands ... 13-48
- show commands, supplicant ... 13-55
- statistics ... 13-48
- supplicant
 - client not using ... 13-31
 - configuring switch port ... 13-46
 - enabling switch port ... 13-46
 - identity option ... 13-46
 - secret ... 13-46
 - switch port operating as ... 13-44
- supplicant state ... 13-55
- supplicant statistics, note ... 13-55
- supplicant, configuring ... 13-44
- supplicant-timeout ... 13-20
- switch username and password ... 13-4
- terminology ... 13-6
- traffic flow on unauthenticated ports ... 13-24
- troubleshooting, gvrp ... 13-56
- trunked port blocked ... 13-13
- tx-period ... 13-20
- unauthenticated port ... 13-24
- unauthorized ... 13-19
- unauthorized-Client VLAN ... 13-21
- unauthorized-client VLAN, defined ... 13-8
- unauthorized-Client VLAN, multiple clients ... 13-36
- unauth-period ... 13-21
- unauth-period command ... 13-30
- unauth-vid ... 13-21
- use model, open VLAN mode ... 13-28
- used with port-security ... 13-42
- user-based
 - access ... 13-4
 - See also* port based
 - authentication ... 13-10
 - client authentication ... 13-4
 - client limit ... 13-3, 13-4, 13-43
 - client-limit, enable ... 13-18
 - clients use same VLAN ... 13-27
 - convert to port-based ... 13-18
 - enable ... 13-17, 13-43
 - limit ... 13-4
 - limit for web auth, MAC auth ... 13-18
 - See also* user-based.
 - tagged VLAN ... 13-4
 - VLAN ... 13-35, 13-36
 - Web/MAC Auth clients ... 13-5
- user-based vs. port-based ... 13-14
- VLAN
 - authorized-client ... 13-32, 13-33, 13-34
 - guest ... 13-33
 - RADIUS assigned, effect ... 13-35
 - RADIUS override ... 13-32
 - RADIUS-assigned ... 13-33
 - tagged ... 13-30, 13-31
 - temporary membership ... 13-33
 - unauthorized-client ... 13-33, 13-34
 - unauthorized-client, best use ... 13-36
 - unauthorized-client, caution ... 13-34
 - unauthorized-client, on different ports ... 13-36

- untagged ... 13-27, 13-30, 13-31
- untagged membership ... 13-18
- VLAN operation ... 13-56
- VLAN use, multiple clients ... 13-6
- VLAN, assignment conflict ... 13-12
- VLAN, membership priority ... 13-10, 13-27
- VLAN, priority, RADIUS ... 13-31
- VLAN, tagged membership ... 13-31
- Wake-on-LAN traffic ... 13-24
- Web/MAC Auth effect ... 13-60

A

- aaa authentication ... 5-8
 - web browser ... 6-11
- aaa port-access
 - See* Web or MAC Authentication.
- access levels, authorized IP managers ... 15-3
- accounting
 - See* RADIUS.
- ACL
 - 802.1X client limit ... 10-20
 - 802.1X, effect on ... 10-20
 - ACE
 - after match not used ... 10-47
 - defined ... 10-10
 - general rules ... 10-49
 - insert in list ... 10-88
 - limit ... 10-33
 - minimum number ... 10-113
 - not used ... 10-29
 - order in list
 - See* sequence, ACEs.
 - ACE, after match not used ... 10-33
 - AppleTalk ... 10-32
 - application methods ... 10-31
 - application point ... 10-24
 - application points ... 10-31
 - applications ... 10-5, 10-6, 10-15, 10-20, 10-24, 10-26, 10-41
 - assign nonexistent i.d. ... 10-48
 - assign to VLAN ... 10-48
 - assigning ... 10-41
 - assigning to a VLAN ... 10-81, 10-82, 10-84
 - assignment not deleted ... 10-85
 - basic structure ... 10-42
 - broadcasts, effect on ... 10-113
 - character limit ... 10-52

- CIDR ... 10-11
 - mask ... 10-50
 - mask bits, IP address ... 10-54, 10-58
- CIDR, mask ... 3-26
- command
 - syntax ... 10-53
- command summary
 - extended ... 10-8
 - standard ... 10-6
- command syntax ... 10-57
- configuration planning ... 10-24
- configured but not used ... 10-48
- configured, not used ... 10-48
- configuring ... 10-41
- configuring offline ... 10-23
- connection-rate ACL ... 10-11, 10-15, 10-16, 10-20
- copy operation appends ... 10-105
- create, CLI method ... 10-49
- DA, defined ... 10-11, 10-13
- defined ... 10-4, 10-10
- definitions ... 10-10
- deleting from config ... 10-85
- deny any, implicit ... 10-23, 10-27, 10-30, 10-31, 10-33, 10-34, 10-43, 10-47, 10-48
- deny any, implicit, supersede ... 10-43
- deny, defined ... 10-11
- destination on the switch ... 10-25
- disable ... 10-9
- display ... 10-9
 - ACLs and assignments ... 10-103
 - assignments ... 10-99, 10-100
 - configuration details ... 10-98
 - content of an ACL ... 10-101
 - data types ... 10-103
 - summary, configured ACLs ... 10-97
- dynamic ... 10-5, 10-6, 10-26
- dynamic port ACL ... 10-11
 - denied traffic ... 10-19
 - multiple clients connected ... 10-19
- dynamic port ACL application ... 10-19
- dynamic port ACL operation defined ... 10-16
- dynamic port joins to a VLAN ... 10-35
- editing ... 10-86
- editing offline ... 10-104
- effect of replacing ... 10-48
- enable ... 10-9
- end ... 10-48
- established ... 10-69

- example, named extended ... 10-73
- exception for connection-rate filtering ... 10-22
- exit statement ... 10-48
- extended
 - command summary ... 10-8
 - configure ... 10-60, 10-74
 - create ... 10-8, 10-60
 - defined ... 10-11, 10-42
 - delete ... 10-9, 10-61
 - named, configure ... 10-62
 - numbered, configure ... 10-75
 - numeric I.D. range ... 10-42
 - protocol options ... 10-42
 - remark ... 10-9, 10-61
 - resequence ... 10-8, 10-60
 - sequence number ... 10-8, 10-60
 - structure ... 10-45
 - use ... 10-15
- features, common to all ... 10-22
- filtering methods ... 10-15
- filtering process ... 10-28, 10-33
- host option ... 10-39
- ICMP
 - code ... 10-79
 - configure ... 10-79
 - options ... 10-70, 10-79
 - traffic ... 10-24
 - type ... 10-79
 - type names ... 10-71
- ID, defined ... 10-11
- identifier ... 10-12
- IGMP
 - configure ... 10-72
 - option ... 10-80
 - traffic ... 10-24
 - type ... 10-80
- implicit deny
 - See deny any, implicit.*
- implicit deny any ... 10-30
- implicit deny, defined ... 10-12
- inbound traffic, defined ... 10-12
- inbound traffic, RACL ... 10-12
- interface assignment, options ... 10-26
- inverse mask
 - See wildcard.*
- ip routing required ... 10-5
- IP routing requirement for RACL ... 10-25
- IPX ... 10-32
- limit ... 10-61
- log function, with mirroring ... 10-23
- log message
 - See ACL, logging.*
- logging ... 10-23, 10-24, 10-55
 - described ... 10-109
 - session ... 10-23
- logging, notes ... 10-113
- mask ... 10-11, 10-23, 10-37, 10-54
 - CIDR ... 10-50
 - defined ... 10-11
 - multiple IP addresses ... 10-40
 - one IP address ... 10-39
- mask, CIDR ... 3-26
- match, always ... 10-48
- match, criteria ... 10-38
- match, example ... 10-39
- match, ignored ... 10-33
- maximum allowed ... 10-33, 10-52
- mirrored traffic exception ... 10-22
- mirroring ... 10-16
- mirroring, log keyword ... 10-16
- mirroring, with log function ... 10-23
- multiple ACLs on interface ... 10-20
- multiple applications ... 10-20
- multiple lists on an interface ... 10-21
- multiple on same interface ... 10-20
- name or number assignment ... 10-48
- name string, maximum characters ... 10-42
- named ... 10-13
- named, character limit ... 10-52
- named, rule ... 10-49
- nonexistent i.d., assign ... 10-48
- non-IP traffic ... 10-32
- number of entries ... 10-22
- numbered ... 10-13
 - manage as named ... 10-52
 - rule ... 10-49
- offline editing ... 10-104
- operator, comparison ... 10-68, 10-69
- outbound traffic, defined ... 10-13
- override implicit deny ... 10-30
- packet screened by multiple lists ... 10-21
- Permit Any forwarding ... 10-13, 10-33
- permit, defined ... 10-13
- permit, with multiple ACLs ... 10-20
- permit/deny policies, defined ... 10-42
- planning ... 10-24, 10-30

- policies ... 10-30
- policy application points ... 1-8, 10-4
- policy type ... 10-42
- policy, permit/deny ... 10-42
- port ... 10-34
- port ACL defined
 - See also* static port ACL and dynamic port ACL ... 10-5
- port ACL operation defined ... 10-16
- port added to trunk ... 10-34
- port removed from trunk ... 10-34
- port-based 802.1X ... 10-20
- port-based security ... 10-20
- ports affected ... 10-35
- precedence ... 10-78
- precedence setting ... 10-24
- precedence, numbers - names ... 10-66
- purpose ... 10-4
- RACL ... 10-12
- RACL application ... 10-16
- RACL configure ... 10-9
- RACL defined ... 10-5
- RACL operation defined ... 10-15
- RACL, screening switched traffic ... 10-21
- RADIUS-assigned ... 10-5, 10-6, 10-19, 10-26
- RADIUS-assigned ACLs ... 10-11
- remark ... 10-13
 - remove from an ACE ... 10-94
- removing from a VLAN ... 10-81, 10-82, 10-84
- replacing ... 10-34
- replacing active ACEs ... 10-48
- resequence ... 10-75
- resource monitor ... 10-114
- routed traffic ... 10-35
- routing requirement ... 10-33
- rules, configuration ... 10-33
- rules, operation ... 10-33
- SA or DA on the switch ... 10-35
- SA, defined ... 10-13
- security use ... 1-8, 10-4, 10-32
- security use, caution ... 10-32
- sequence number ... 10-14, 10-87
 - use to delete ACE ... 10-90
 - use to insert ACE ... 10-88
- sequence number interval ... 10-75
- source routing, caution ... 10-25, 10-41
- standard
 - command summary ... 10-6
 - configure ... 10-51
 - create ... 10-6, 10-51
 - defined ... 10-14, 10-42
 - delete ... 10-51
 - example ... 10-59
 - named, configure ... 10-54
 - numbered, configure ... 10-57
 - numeric I.D. range ... 10-42
 - remark ... 10-6, 10-51
 - resequence ... 10-6, 10-51
 - sequence number ... 10-6, 10-51
 - structure ... 10-43
 - use ... 10-15, 10-52
- static port ACL ... 10-9
- static port ACL application ... 10-19
- static VLAN requirement ... 10-25, 10-34, 10-35
- static, defined ... 10-5
- static-port ACL ... 10-14
- supernetting ... 10-37
- supersede implicit deny any ... 10-47
- switched packets ... 10-35
- syntax
 - See* command syntax.
- Syslog
 - See* ACL, logging.
- TCP or UDP port number, IANA ... 10-69
- TCP/UDP operators ... 10-68
- TCP/UDP, port names ... 10-69
- terms ... 10-10
- ToS setting ... 10-24
- ToS, numbers - names ... 10-67, 10-78
- traffic not filtered ... 10-32
- traffic to/from the switch ... 10-35
- traffic types filtered ... 10-4, 10-30
- trunk ... 10-34
- trunk, adding port ... 10-34
- type ... 10-47, 10-52, 10-89, 10-96, 10-99, 10-100, 10-101
- user-based 802.1X ... 10-20
- user-based security ... 10-20
- VACL application ... 10-18
- VACL configure ... 10-9
- VACL defined ... 10-5
- VACL operation defined ... 10-16
- VLANs ... 10-34
- where applied to traffic ... 10-26, 10-35
- wildcard ... 10-11, 10-38, 10-39
- wildcard, defined ... 10-14

- ACL, connection-rate
 - See* connection-rate filtering
- ACLs
 - management access protection ... 1-8
 - See also* RADIUS-assigned ACLs.
- address
 - authorized for port security ... 14-5
- alerts
 - generating for monitored events ... 11-23
- ARP
 - adding IP-to-MAC binding ... 11-18
 - debugging ... 11-21
 - displaying statistics ... 11-21
 - dynamic ARP protection ... 11-15
 - requests ... 11-22
 - trusted ports ... 11-17
 - validation checks on ARP packets,
 - configuring ... 11-19
- authentication failures ... 11-23
- authentication, RADIUS override ... 7-4
- authenticator
 - backend state ... 13-51
 - state ... 13-51
- authorized addresses
 - for IP management security ... 15-4
 - for port security ... 14-5
- authorized IP managers
 - access levels ... 15-3
 - building IP masks ... 15-9
 - configuring in browser interface ... 15-7, 15-9
 - configuring in console ... 15-5
 - definitions of single and multiple ... 15-4
 - effect of duplicate IP addresses ... 15-12
 - IP mask for multiple stations ... 15-10
 - IP mask for single station ... 15-9
 - IP mask operation ... 15-4
 - operating notes ... 15-12
 - overview ... 15-1
 - troubleshooting ... 15-12
- authorized server ... 11-4
- authorized server address, configuring ... 11-8

B

- binding database ... 11-11
- BPDU filtering ... 1-12
- BPDU protection ... 1-12
- broadcast traffic

- effect of ACL ... 10-113

C

- certificate
 - CA-signed ... 9-3
 - root ... 9-4
 - self-signed ... 9-3
- Class of Service ... 7-3, 7-5, 7-6
 - RADIUS ... 7-4
- Clear button ... 1-6
 - to delete password protection ... 2-6
- configuration
 - filters ... 12-2
 - port security ... 14-7
 - RADIUS
 - See RADIUS.*
 - SSH
 - See SSH.*
- connection-rate ACL ... 3-6
- connection-rate filtering ... 3-31
 - access-control list ... 3-6
- ACL
 - ACE mask ... 3-26
 - application to port ... 3-21
 - applying ... 3-26
 - CIDR notation ... 3-26
 - configuring ... 3-19
 - example ... 3-27, 3-28
 - filter ... 3-20, 3-22, 3-23, 3-30
 - ignore ... 3-20, 3-22, 3-23, 3-29
 - implicit filter ... 3-20, 3-30
 - operation ... 3-20
 - source IP ... 3-21
 - UDP/TCP ... 3-23
 - UDP/TCP options ... 3-24
 - UDP/TCP port data ... 3-25
 - VLAN effect ... 3-20
- activation ... 3-4, 3-7
- benefits ... 3-4
- block ... 3-12
- blocked host ... 3-5, 3-7, 3-18
- blocked host, re-enable ... 3-5, 3-9
- configuration, example ... 3-14
- configuring per-port ... 3-12
- edge of network ... 3-3
- enabling, commands ... 3-11
- event log notice ... 3-5

- false positive ... 3-6
- guidelines ... 3-8, 3-9
- high rate, legitimate ... 3-18
- host, trusted ... 3-18
- host, unblocking ... 3-18
- ICMP ping message ... 3-3
- notify and reduce ... 3-5
- notify only ... 3-5
- notify-only ... 3-12
- operating rules ... 3-7
- operation ... 3-5
- options ... 3-5
- penalty period, throttling ... 3-12
- port setting change, effect ... 3-7
- reboot, effect ... 3-7
- recommended application ... 3-3
- re-enable blocked host ... 3-7
- routed traffic ... 3-10
- sensitivity level ... 3-5, 3-8
- sensitivity level, changing ... 3-18
- sensitivity level, command ... 3-11
- show, command ... 3-15, 3-17
- signature recognition ... 3-3, 3-4
- SNMP trap ... 3-5
- throttle ... 3-5, 3-6, 3-12
- trigger ... 3-4, 3-7, 3-10
- unblock command ... 3-9, 3-18
- unblocking a host ... 3-7
- VLAN delete, effect ... 3-7
- worm ... 3-3, 3-4
- connection-rate filtering, exception ... 10-22
- console, for configuring
 - authorized IP managers ... 15-5
- CoS ... 7-3, 7-4, 7-5, 7-6
 - priority assignment ... 6-4
 - RADIUS override ... 7-5
- CoS override ... 13-48, 13-51

D

- DA, defined ... 7-9, 10-11, 10-13
- database, snooping ... 11-4
- debug logging, DHCP snooping ... 11-12
- default configuration and security ... 1-3
- denial-of-service (DoS) attack ... 1-12, 11-3, 11-22
- DES ... 8-3, 9-3
- DHCP Option 82
 - IP-to-MAC binding database ... 11-18

- DHCP protection
 - See* DHCP snooping.
- DHCP snooping ... 11-3
 - authorized server ... 11-4
 - binding database ... 11-11
 - changing remote-id ... 11-10
 - configuring authorized server address ... 11-8
 - database ... 11-4
 - denial-of-service attack ... 11-3
 - disable MAC check ... 11-10
 - disabling ... 11-4
 - dropping packets ... 11-4
 - enabling ... 11-4
 - debug logging ... 11-12
 - on trusted ports ... 11-7
 - on VLANs ... 11-5, 11-6
 - IP-to-MAC binding database ... 11-18
 - log messages ... 11-13
 - Option 82 ... 11-8
 - option parameter ... 11-5
 - remote-id ... 11-9
 - show configuration ... 11-5
 - stats ... 11-5
 - trust ... 11-5
 - untrusted-policy ... 11-9
 - verify ... 11-5
- DHCPACK ... 11-4
- DHCPDECLINE ... 11-4
- DHCPNACK ... 11-4
- DHCPOFFER ... 11-4
- DHCPRELEASE ... 11-4
- duplicate IP address
 - effect on authorized IP managers ... 15-12
- dynamic ARP protection
 - additional validation checks on ARP packets ... 11-19
 - ARP packet debugging ... 11-21
 - displaying ARP statistics ... 11-21
 - IP-to-MAC binding, adding to DHCP database ... 11-18
 - trusted ports, configuring ... 11-17
 - verifying configuration ... 11-20
- dynamic ARP protection, enabling ... 11-15

E

- Eavesdrop Protection ... 14-4
- enhancing network security ... 7-12

- event log
 - alerts for monitored events ... 11-23
 - connection-rate filtering alerts ... 3-31
 - intrusion alerts ... 14-40
 - messages ... 3-31

F

- filter, source-port
 - applicable models ... 12-2
 - editing ... 12-20
 - filter indexing ... 12-22
 - filter type ... 12-8
 - idx ... 12-8, 12-22
 - index ... 12-8, 12-22
 - operating rules ... 12-4, 12-6
 - port-trunk operation ... 12-3, 12-19
 - show ... 12-8
 - value ... 12-8
 - viewing ... 12-8
- filters ... 12-2
 - effect of IGMP ... 12-16
 - multicast ... 12-15
 - protocol ... 12-16
 - source port ... 12-4
 - source-port filter value ... 12-22
 - static ... 12-3
 - types ... 12-3
- front-panel access ... 1-6

G

- guest VLAN ... 13-7, 13-8, 13-27
- GVRP ... 13-56
- GVRP, static VLAN not advertised ... 13-59

I

- IANA ... 10-69
- IANA, protocol numbers ... 10-65, 10-70
- Identity Driven Manager
 - See* IDM.
- IDM ... 1-13, 7-2, 7-8, 7-29
 - See also* RADIUS-assigned ACLs
 - RADIUS-assigned ACLs.
- IGMP
 - effect on filters ... 12-16
 - IP multicast address range ... 12-16

- inconsistent value, message ... 14-21
- intrusion alarms
 - entries dropped from log ... 14-42
 - event log ... 14-40
 - prior to ... 14-42

Intrusion Log

- prior to ... 14-38, 14-39

IP

- address count ... 11-22
- authorized IP managers ... 15-1
- reserved port numbers ... 8-17

IP masks

- building ... 15-9
- for multiple authorized manager stations ... 15-10
- for single authorized manager station ... 15-9
- operation ... 15-4

IP routing

- dynamic ARP protection, enabling ... 11-15
- required for ACLs ... 10-5
- validation checks on ARP packets, configuring ... 11-19

IP-to-MAC binding ... 11-18

K

key chain

- See* KMS key chain.

key management system

- See* KMS.

KMS

- accept key time ... 16-5, 16-7
- assigning a time-dependent key ... 16-5
- assigning a time-independent key ... 16-4
- generating a key chain ... 16-3
- generating a time-dependent key ... 16-5
- generating a time-independent key ... 16-4
- key chain ... 16-2
- key chain entry ... 16-3
- key chain generation ... 16-3
- overview ... 1-11, 16-2
- send key time ... 16-5
- time protocol ... 16-6
- time-dependent key ... 16-2, 16-5, 16-6
- time-independent key ... 16-2, 16-4

L

LACP

802.1X not allowed ... 13-13, 13-17, 13-61

log keyword, ACL mirroring ... 10-16

login attempts, monitoring ... 11-23

M

MAC addresses

monitoring activity ... 11-23

MAC auth

port access ... 13-4

MAC Authentication

authenticator operation ... 4-5

blocked traffic ... 4-4

CHAP

defined ... 4-9

usage ... 4-4

client status ... 4-39

configuration commands ... 4-25, 4-31

configuring

access control on unauthenticated
ports ... 4-34

controlled directions ... 4-34

on the switch ... 4-24, 4-31

switch for RADIUS access ... 4-14

the RADIUS server ... 4-13

features ... 4-3

general setup ... 4-12

LACP not allowed ... 4-12

rules of operation ... 4-10

show status and configuration ... 4-36

terminology ... 4-9

Wake-on-LAN traffic ... 4-35

MAC authentication

overview ... 1-9

MAC Lockdown ... 14-3

MAC Lockout ... 14-3

manager password ... 1-3, 2-3, 2-5, 2-6

manager password recommended ... 5-7

MD5

See RADIUS.

message

inconsistent value ... 14-21

MIB (Management Information Base)

SNMP access ... 1-4

SNMP access to authentication MIB ... 1-5

mirrored traffic ... 10-22

mirroring ... 10-16

multicast address, spanning tree protocol ... 12-16

multicast filter ... 12-3, 12-15

multicast MAC address, STP

N

named source port filters

configuring ... 12-7

operating rules ... 12-6

viewing ... 12-8

NAS ... 7-10

network management applications ... 7-2

O

open VLAN mode

See 802.1X access control.

OpenSSH ... 8-3

OpenSSL ... 9-2

operating notes

authorized IP managers ... 15-12

port security ... 14-42

operator password ... 2-3, 2-5, 2-6

Option 82

snooping ... 11-5

P

packet validation ... 11-5

password

browser/console access ... 2-4

case-sensitive ... 2-5

caution ... 2-4

default configuration ... 1-3

delete ... 2-6

deleting with the Clear button ... 2-6

if you lose the password ... 2-6

incorrect ... 2-4

length ... 2-5

operator only, caution ... 2-4

pair ... 2-2

protection ... 1-3

setting ... 2-5

SNMP configuration ... 2-2, 2-8

password pair ... 2-2

password security ... 1-3, 8-18

PCM

- See* ProCurve Manager.
- physical security ... 1-6
- port
 - security configuration ... 14-3
 - trusted ... 11-17
 - untrusted ... 11-18
- port access
 - client limit ... 13-18
 - concurrent ... 13-18
 - MAC auth ... 13-4
 - See also* 802.1X access control.
 - tracking client authentication failures ... 11-23
 - Web auth ... 13-4
 - Web/MAC ... 13-18
- port ACL ... 10-5
- port monitoring, ACL ... 10-16
- port scan, detecting ... 11-22
- port security
 - 802.1X, learn mode requirement ... 14-14
 - authorized address definition ... 14-5
 - basic operation ... 14-4
 - caution, device limit ... 14-14
 - configuring ... 14-7
 - configuring in browser interface ... 14-34, 14-41
 - event log ... 14-40
 - notice of security violations ... 14-34
 - operating notes ... 14-42
 - overview ... 1-10, 14-3
 - prior to ... 14-42
 - proxy web server ... 14-42
 - TCP/UDP closed ports ... 11-22
- port-based access control
 - port-security learn mode ... 14-14
 - Rate-Limit override ... 13-48
 - See* 802.1X access control.
 - VLAN, tagged member ... 13-28
- ports
 - trusted ... 11-5
- prior to ... 14-38, 14-39, 14-42
- Privacy Enhanced Mode (PEM)
 - See* SSH.
- ProCurve Manager ... 1-13, 7-2, 14-4
- protocol filters ... 12-16
- proxy
 - web server ... 14-42

R

- RACL defined ... 10-5
- RADIUS
 - accounting ... 6-4, 6-32
 - accounting, configuration outline ... 6-34
 - accounting, configure server access ... 6-35
 - accounting, configure types on switch ... 6-36
 - accounting, exec ... 6-33, 6-36
 - accounting, interim updating ... 6-38
 - accounting, network ... 6-36, 6-37
 - accounting, operating rules ... 6-33
 - accounting, server failure ... 6-34
 - accounting, session-blocking ... 6-38
 - accounting, start-stop method ... 6-37
 - accounting, statistics terms ... 6-41
 - accounting, stop-only method ... 6-37
 - accounting, system ... 6-33, 6-36
 - ACL, dynamic port ... 7-15
 - administrative-user service-type value ... 6-12
 - authentication options ... 6-3
 - authentication, local ... 6-22
 - authentication, web ... 6-3, 6-10
 - authentication, web browser ... 6-11, 6-13
 - authorization ... 6-24
 - bypass RADIUS server ... 6-11
 - Class of Service ... 7-3, 7-5, 7-6
 - commands accounting ... 6-33
 - commands authorization ... 6-24
 - commands, accounting ... 6-32
 - commands, switch ... 6-8
 - configuration outline ... 6-9
 - configure server access ... 6-13
 - configuring server ... 6-26
 - configuring switch global parameters ... 6-15
 - CoS override ... 7-3
 - dynamic port ACL ... 7-9, 7-11
 - general setup ... 6-7
 - HP-Command-Exception ... 6-26
 - HP-command-string ... 6-26
 - local authentication ... 6-11
 - login privilege-mode, application options ... 6-12
 - login-privilege mode ... 6-12
 - manager access denied ... 6-12
 - manager access privilege ... 6-12
 - manager access privilege, service type
 - value ... 6-8
 - MD5 ... 6-6
 - messages ... 6-47

- multiple ACL application types in use ... 7-15
 - NAS-Prompt-User service-type value ... 6-12
 - network accounting ... 6-32
 - operating rules, switch ... 6-6
 - override CoS ... 7-5
 - override CoS, example ... 7-5, 7-6
 - override Rate-Limiting ... 7-5
 - override Rate-Limiting, example ... 7-5, 7-6
 - override, precedence, multiple clients ... 7-7
 - rate-limiting ... 7-3, 7-4, 7-6
 - Rate-Limiting override ... 7-3
 - security ... 6-11
 - security note ... 6-4
 - server access order ... 6-33
 - server access order, changing ... 6-44
 - servers, multiple ... 6-17
 - service type value ... 6-8
 - service-type value ... 6-12
 - service-type value, null ... 6-12
 - show accounting ... 6-43
 - show authentication ... 6-42
 - SNMP access security not supported ... 6-4
 - SNMP access to auth config MIB ... 6-4
 - statistics, viewing ... 6-40
 - terminology ... 6-5
 - TLS ... 6-6
 - vendor specific attributes ... 6-26
 - vendor-specific attributes ... 7-3
 - VSAAs ... 6-27
 - web browser security not supported ... 6-7
 - web-browser access controls ... 6-23
 - web-browser security not supported ... 6-4, 6-23
 - RADIUS-assigned ACLs ... 7-8, 10-5
 - 802.1X port-based access ... 7-18
 - 802.1X user-based access ... 7-18
 - ACE, defined ... 7-8
 - application type ... 7-9
 - contrasting dynamic and static ... 7-13
 - DA, defined ... 7-9
 - defined ... 7-8
 - definitions ... 7-8
 - deny any, implicit, switched packets ... 7-17
 - deny in any ACL on an interface ... 7-18
 - deny, defined ... 7-9
 - dynamic port ... 7-11, 7-15
 - dynamic port ACL ... 7-9
 - dynamic port ACL, effect ... 7-18
 - filters ... 7-12
 - implicit deny, defined ... 7-9
 - inbound traffic, defined ... 7-10
 - inverse mask
 - See* wildcard.
 - mask ... 7-9
 - mask, defined ... 7-9
 - multiple application types in use ... 7-15
 - multiple clients, access restriction ... 7-18
 - multiple dynamic ACLs ... 7-18
 - multiple, on an interface ... 7-17
 - outbound traffic, defined ... 7-10
 - permit, defined ... 7-10
 - RADIUS-based ... 7-15
 - resource monitor ... 7-29
 - See also* ACLs.
 - source routing, caution ... 7-14
 - static-port ACL ... 7-10
 - switched packets ... 7-17
 - terminology ... 7-8
 - terms ... 7-8
 - wildcard ... 7-9, 7-10
 - wildcard, defined ... 7-10
 - RADIUS-based ACL filtering ... 7-15
 - Rate-Limit override ... 13-48, 13-51
 - rate-limiting ... 6-4, 7-3, 7-4, 7-6, 11-24
 - Rate-Limiting, RADIUS override ... 7-5
 - remote access security ... 1-4
 - remote-id
 - changing ... 11-10
 - snooping ... 11-9
 - reserved port numbers ... 8-17, 9-20
 - Reset button ... 1-6
 - resource monitor
 - See Management and Configuration Guide.*
 - routing
 - source-routing, caution ... 7-14, 10-25, 10-41
- S**
- SA ... 10-13
 - secure copy ... 1-6
 - secure file transfers ... 1-6
 - secure management VLAN ... 1-7
 - security
 - authorized IP managers ... 15-1
 - per port ... 14-3
 - security violations
 - detecting ... 11-22

- notices of ... 14-34
- security, ACL
 - See* ACL, security use.
- security, password
 - See* SSH.
- setting a password ... 2-5
- SFTP ... 1-6
- SNMP
 - authentication failures ... 11-23
 - disabling access to authentication MIB ... 1-5
 - password and username configuration ... 2-2, 2-8
 - RADIUS access to auth config MIB ... 6-4
 - trap generation ... 11-25
- snooping
 - authorized server ... 11-4
 - authorized server address ... 11-8
 - binding database ... 11-11
 - changing remote-id ... 11-10
 - DHCP ... 11-3
 - disable MAC check ... 11-10
 - Option 82 ... 11-5, 11-8
 - statistics ... 11-5
 - untrusted-policy ... 11-9
 - verify ... 11-5
- source port filter ... 12-3
- source port filters ... 12-4
 - named ... 12-6
- source-routing, caution ... 7-14, 10-25, 10-41
- spanning tree
 - caution about filtering ... 12-16
 - edge port configuration ... 4-22, 4-35, 13-24
 - security features ... 1-12
- spanning tree protocol
 - See* STP.
- SSH
 - authenticating switch to client ... 8-3
 - authentication, client public key ... 8-2
 - authentication, user password ... 8-2
 - caution, security ... 8-18
 - CLI commands ... 8-9
 - client behavior ... 8-15, 8-16
 - client public-key authentication ... 8-19, 8-22
 - client public-key, clearing ... 8-26
 - client public-key, creating file ... 8-23
 - client public-key, displaying ... 8-25
 - configuring authentication ... 8-18
 - crypto key ... 8-11
 - disabling ... 8-11
 - enable ... 8-16, 9-19
 - enabling ... 8-15
 - erase host key pair ... 8-11
 - generate host key pair ... 8-11
 - generating key pairs ... 8-10
 - host key pair ... 8-11
 - key, babble ... 8-11
 - key, fingerprint ... 8-11
 - keys, zeroing ... 8-11
 - key-size ... 8-17
 - known-host file ... 8-13, 8-15
 - man-in-the-middle spoofing ... 8-16
 - messages, operating ... 8-27
 - OpenSSH ... 8-3
 - operating rules ... 8-8
 - outbound SSH not secure ... 8-8
 - overview ... 1-9
 - password security ... 8-18
 - password-only authentication ... 8-18
 - passwords, assigning ... 8-9
 - PEM ... 8-4
 - prerequisites ... 8-5
 - public key ... 8-5, 8-13
 - public key, displaying ... 8-14
 - reserved IP port numbers ... 8-17
 - security ... 8-18
 - SSHv2 ... 8-2
 - steps for configuring ... 8-6
 - supported encryption methods ... 8-3
 - switch key to client ... 8-12
 - terminology ... 8-3, 16-2
 - unauthorized access ... 8-27
 - version ... 8-2
 - zeroing a key ... 8-11
 - zeroize ... 8-11
- SSL
 - CA-signed ... 9-3, 9-15
 - CA-signed certificate ... 9-3, 9-15
 - CLI commands ... 9-7
 - client behavior ... 9-17, 9-18
 - crypto key ... 9-10
 - disabling ... 9-10, 9-17
 - enabling ... 9-17
 - erase certificate key pair ... 9-10
 - erase host key pair ... 9-10
 - generate CA-signed ... 9-15
 - generate CA-signed certificate ... 9-15

- generate host key pair ... 9-10
- generate self-signed ... 9-13
- generate self-signed certificate ... 9-10, 9-13
- generate server host certificate ... 9-10
- generating Host Certificate ... 9-9
- host key pair ... 9-10
- key, babble ... 9-12
- key, fingerprint ... 9-12
- man-in-the-middle spoofing ... 9-18
- OpenSSL ... 9-2
- operating notes ... 9-6
- operating rules ... 9-6
- passwords, assigning ... 9-7
- prerequisites ... 9-5
- remove self-signed certificate ... 9-10
- remove server host certificate ... 9-10
- reserved TCP port numbers ... 9-20
- root ... 9-4
- root certificate ... 9-4
- self-signed ... 9-3, 9-13
- self-signed certificate ... 9-3, 9-10, 9-13
- server host certificate ... 9-10
- SSL server ... 9-3
- SSLv3 ... 9-2
- steps for configuring ... 9-5
- supported encryption methods ... 9-3
- terminology ... 9-3
- TLSv1 ... 9-2
- troubleshooting, operating ... 9-21
- unsecured web browser access ... 9-18
- version ... 9-2
- zeroize ... 9-10
- static ACL defined ... 10-5
- static filter limit ... 12-3
- static multicast filter ... 12-15
- STP
 - prerequisite for 802.1X controlled directions ... 13-24
 - prerequisite for MAC-based controlled directions ... 4-35
 - prerequisite for web-based controlled directions ... 4-22
 - STP multicast MAC address
- supernetting ... 10-37
- supersede implicit deny any ... 10-43
- Syslog
 - See* ACL, logging.
- system delay ... 11-23

- system resource usage ... 11-22

T

TACACS

- aaa parameters ... 5-12
- authentication ... 5-3
- authentication process ... 5-20
- authentication, local ... 5-22
- authorized IP managers, effect ... 5-25
- configuration, authentication ... 5-11
- configuration, encryption key ... 5-19
- configuration, server access ... 5-15
- configuration, timeout ... 5-20
- configuration, viewing ... 5-10
- encryption key ... 5-6, 5-15, 5-16, 5-19
- encryption key exclusion ... 5-26
- encryption key, general operation ... 5-23
- encryption key, global ... 5-20
- general operation ... 5-2
- IP address, server ... 5-15
- local manager password requirement ... 5-26
- messages ... 5-25
- NAS ... 5-3
- precautions ... 5-5
- preparing to configure ... 5-8
- preventing switch lockout ... 5-15
- privilege level code ... 5-7
- server access ... 5-15
- server priority ... 5-18
- setup, general ... 5-5
- show authentication ... 5-8
- system requirements ... 5-5
- TACACS+ server ... 5-3
- testing ... 5-5
- TFTP, configuration ... 5-26
- timeout ... 5-15
- troubleshooting ... 5-6
- unauthorized access, preventing ... 5-7
- web access, controlling ... 5-24
- web access, no effect on ... 5-5
- tacacs-server ... 5-8
- TCP
 - reserved port numbers ... 9-20
- TCP/UDP
 - monitoring packets to closed ports ... 11-22
- Telnet ... 1-4
- test ... 5-15

TLS

See RADIUS.

troubleshooting

authentication via Telnet ... 5-15
authorized IP managers ... 15-12

trunk

filter, source-port ... 12-3, 12-19
LACP, 802.1X not allowed ... 13-17
port added or removed, ACL ... 10-34
See also LACP.

trusted port

defined ... 11-17
enabling ... 11-17

trusted ports ... 11-5

enabling snooping ... 11-7

U

untrusted policy, snooping ... 11-9

user name

cleared ... 2-6
SNMP configuration ... 2-2, 2-8

V

VACL defined ... 10-5

value, inconsistent ... 14-21

vendor specific attributes ... 6-27

Vendor-Specific Attribute ... 7-10

vendor-specific attribute

configuring ... 7-3

vendor-specific attributes ... 7-3

virus detection

monitoring ARP requests ... 11-22

virus-throttling

See connection-rate filtering.

VLAN

802.1X ... 13-56

802.1X, ID changes ... 13-59

802.1X, suspend untagged VLAN ... 13-52

connection-rate filtering ... 3-20

not advertised for GVRP ... 13-59

secure management ... 1-7

VSA ... 7-10

See vendor-specific attribute.

VSAs, defining ... 6-28

W

Wake-on-LAN

on 802.1X-aware ports ... 13-24
on MAC-authenticated ports ... 4-35
on web-authenticated ports ... 4-23

warranty ... 1-ii

Web auth

port access ... 13-4

Web Authentication

authenticator operation ... 4-5

blocked traffic ... 4-4

CHAP

defined ... 4-9

usage ... 4-4

client status ... 4-39

configuration commands ... 4-18

configuring

access control on unauthenticated
ports ... 4-22

controlled directions ... 4-22

on the switch ... 4-17

switch for RADIUS access ... 4-14

features ... 4-3

general setup ... 4-12

LACP not allowed ... 4-12

redirect URL ... 4-9

rules of operation ... 4-10

show status and configuration ... 4-28

terminology ... 4-9

Wake-on-LAN traffic ... 4-23

Web authentication

aaa authentication ... 6-8

overview ... 1-9

Web browser access ... 1-4

Web browser authentication ... 6-8

web browser interface

configuring

port security ... 14-41

configuring port security ... 14-34

SSL ... 9-18

unsecured access, SSL ... 9-18

web browser interface, for configuring

authorized IP managers ... 15-7, 15-9

web server, proxy ... 14-42

Webui access ... 6-7

wildcard

See ACL.

wildcard, ACL, defined ... 7-10, 10-14



Technical information in this document
is subject to change without notice.

© Copyright 2005-2007
Hewlett-Packard Development Company, L.P.

Reproduction, adaptation, or translation
without prior written permission is prohibited
except as allowed under the copyright laws.

February 2007
Manual Part Number
5991-3828