



HP ProCurve Switch Software

Management and Configuration Guide

3500 switches
3500yl switches
5400zl switches
6200yl switches
6600 switches
8200zl switches

Software version K.14.34
September 2009

HP ProCurve
3500 Switches
3500yl Switches
5400zl Switches
6200yl Switch
6600 Switches
8200zl Switches

September 2009
K.14.34

Management and Configuration
Guide

© Copyright 2005–2009 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

5992-3059
September 2009

Applicable Products

HP ProCurve Switch 3500-24	(J9470A)
HP ProCurve Switch 3500-48	(J9472A)
HP ProCurve Switch 3500-24-PoE	(J9471A)
HP ProCurve Switch 3500-48-PoE	(J9473A)
HP ProCurve Switch 3500yl-24G-PWR	(J8692A)
HP ProCurve Switch 3500yl-48G-PWR	(J8693A)
HP ProCurve Switch 5406zl	(J8697A)
HP ProCurve Switch 5406zl-48G-PoE+	(J9447A)
HP ProCurve Switch 5412zl	(J8698A)
HP ProCurve Switch 5412zl-96G-PoE+	(J9448A)
HP ProCurve Switch 6200yl-24G	(J8992A)
HP ProCurve Switch 8206zl	(J9475A)
HP ProCurve Switch 8212zl	(J8715A/B)
HP ProCurve Switch 6600-24G	(J9263A)
HP ProCurve Switch 6600-24G-4XG	(J9264A)
HP ProCurve Switch 6600-24G-24XG	(J9265A)
HP ProCurve Switch 6600-48G	(J9451A)
HP ProCurve Switch 6600-48G-4XG	(J9452A)

HP ProCurve 24-Port 10/100/1000 PoE+ zl Module	(J9307A)
HP ProCurve 20-Port 10/100/1000 PoE+/4-Port MiniGBIC zl Module	(J9308A)
HP ProCurve 4-Port 10GbE SFP+ zl Module	(J9309A)
HP ProCurve 24-Port 10/100 PoE+ zl Module	(J9478A)

Trademark Credits

Microsoft, Windows, and Microsoft Windows NT are US registered trademarks of Microsoft Corporation. Java™ is a US trademark of Sun Microsystems, Inc.

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	xxv
Printed Publications	xxv
Electronic Publications	xxv
Software Feature Index	xxvi

1 Getting Started

Contents	1-1
Introduction	1-2
Conventions	1-2
Command Syntax Statements	1-2
Command Prompts	1-3
Screen Simulations	1-3
Configuration and Operation Examples	1-3
Keys	1-3
Sources for More Information	1-4
Getting Documentation From the Web	1-6
Online Help	1-6
Menu Interface	1-6
Command Line Interface	1-7
Web Browser Interface	1-7
Need Only a Quick Start?	1-8
IP Addressing	1-8
To Set Up and Install the Switch in Your Network	1-8
Physical Installation	1-8

2 Selecting a Management Interface

Contents	2-1
Overview	2-2
Understanding Management Interfaces	2-2
Advantages of Using the Menu Interface	2-3
Advantages of Using the CLI	2-4
General Benefits	2-4
Information on Using the CLI	2-4
Advantages of Using the Web Browser Interface	2-5
Advantages of Using ProCurve Manager or ProCurve Manager Plus	2-7
Custom Login Banners for the Console and Web Browser Interfaces	2-9
Banner Operation with Telnet, Serial, or SSHv2 Access	2-9
Banner Operation with Web Browser Access	2-9
Configuring and Displaying a Non-Default Banner	2-10
Example of Configuring and Displaying a Banner	2-11
Operating Notes	2-13

3 Using the Menu Interface

Contents	3-1
Overview	3-2
Starting and Ending a Menu Session	3-3
How To Start a Menu Interface Session	3-4
How To End a Menu Session and Exit from the Console:	3-5
Main Menu Features	3-7
Screen Structure and Navigation	3-9
Rebooting the Switch	3-12
Menu Features List	3-14
Where To Go From Here	3-15

4 Using the Command Line Interface (CLI)

Contents	4-1
Overview	4-2
Accessing the CLI	4-2
Using the CLI	4-2
Privilege Levels at Logon	4-3
Privilege Level Operation	4-4
Operator Privileges	4-4
Manager Privileges	4-5
How To Move Between Levels	4-7
Listing Commands and Command Options	4-8
Listing Commands Available at Any Privilege Level	4-8
Listing Command Options	4-10
Displaying CLI “Help”	4-11
Configuration Commands and the Context Configuration Modes ..	4-13
CLI Control and Editing	4-16
Executing a Prior Command—Redo	4-16
Repeating Execution of a Command	4-16
Using a Command Alias	4-18
CLI Shortcut Keystrokes	4-20

5 Using the ProCurve Web Browser Interface

Contents	5-1
Overview	5-3
General Features	5-4
Starting a Web Browser	
Interface Session with the Switch	5-5
Using a Standalone Web Browser in a PC or UNIX Workstation	5-5
Using ProCurve Manager (PCM) or ProCurve Manager Plus (PCM+)	5-6
Tasks for Your First ProCurve Web Browser Interface Session ..	5-8
Viewing the “First Time Install” Window	5-8
Security: Creating Usernames and Passwords in the Browser Interface	5-9

Entering a User Name and Password	5-11
Using a User Name	5-11
If You Lose the Password	5-11
Online Help for the Web Browser Interface	5-12
Support/Mgmt URLs Feature	5-13
Support URL	5-14
Help and the Management Server URL	5-14
Using the PCM Server for Switch Web Help	5-15
Status Reporting Features	5-17
The Overview Window	5-17
The Port Utilization and Status Displays	5-18
Port Utilization	5-18
Port Status	5-20
The Alert Log	5-21
Sorting the Alert Log Entries	5-21
Alert Types and Detailed Views	5-22
Status Indicators	5-23
Setting Fault Detection Policy	5-24

6 Switch Memory and Configuration

Contents	6-1
Overview	6-3
Configuration File Management	6-3
Using the CLI To Implement Configuration Changes	6-6
Using the Menu and Web Browser Interfaces To Implement Configuration Changes	6-10
Menu: Implementing Configuration Changes	6-10
Using Save and Cancel in the Menu Interface	6-10
Rebooting from the Menu Interface	6-11
Web: Implementing Configuration Changes	6-13
Using Primary and Secondary Flash Image Options	6-14
Displaying the Current Flash Image Data	6-14
Switch Software Downloads	6-16
Local Switch Software Replacement and Removal	6-17

Rebooting the Switch	6-19
Operating Notes about Booting	6-19
Boot and Reload Command Comparison	6-20
Setting the Default Flash	6-21
Booting from the Default Flash (Primary or Secondary)	6-22
Booting from a Specified Flash	6-23
Using Reload	6-24
Multiple Configuration Files	6-26
General Operation	6-27
Transitioning to Multiple Configuration Files	6-29
Listing and Displaying Startup-Config Files	6-30
Viewing the Startup-Config File Status with Multiple Configuration Enabled	6-30
Displaying the Content of A Specific Startup-Config File	6-31
Changing or Overriding the Reboot Configuration Policy	6-31
Managing Startup-Config Files in the Switch	6-33
Renaming an Existing Startup-Config File	6-34
Creating a New Startup-Config File	6-34
Erasing a Startup-Config File	6-35
Using the Clear + Reset Button Combination To Reset the Switch to Its Default Configuration	6-37
Transferring Startup-Config Files To or From a Remote Server	6-38
TFTP: Copying a Configuration File to a Remote Host	6-38
TFTP: Copying a Configuration File from a Remote Host	6-39
Xmodem: Copying a Configuration File to a Serially Connected Host	6-40
Xmodem: Copying a Configuration from a Serially Connected Host	6-40
Operating Notes for Multiple Configuration Files	6-40
Automatic Configuration Update with DHCP Option 66	6-41
CLI Command	6-41
Possible Scenarios for Updating the Configuration File	6-42
Operating Notes	6-42
Log Messages	6-43

7 Interface Access and System Information

Contents	7-1
Overview	7-2
Interface Access: Console/Serial Link, Web, and Inbound Telnet .	7-3
Menu: Modifying the Interface Access	7-4
CLI: Modifying the Interface Access	7-5
Denying Interface Access by Terminating Remote Management Sessions	7-11
System Information	7-12
Menu: Viewing and Configuring System Information	7-13
CLI: Viewing and Configuring System Information	7-14
Web: Configuring System Parameters	7-18

8 Configuring IP Addressing

Contents	8-1
Overview	8-2
IP Configuration	8-2
Just Want a Quick Start with IP Addressing?	8-4
IP Addressing with Multiple VLANs	8-4
Menu: Configuring IP Address, Gateway, and Time-To-Live (TTL) . .	8-5
CLI: Configuring IP Address, Gateway, and Time-To-Live (TTL)	8-7
Web: Configuring IP Addressing	8-11
How IP Addressing Affects Switch Operation	8-11
DHCP/Bootp Operation	8-12
Network Preparations for Configuring DHCP/Bootp	8-15
Loopback Interfaces	8-16
Introduction	8-16
Configuring a Loopback Interface	8-17
Displaying Loopback Interface Configurations	8-18
IP Preserve: Retaining VLAN-1 IP Addressing Across Configuration File Downloads	8-21
Operating Rules for IP Preserve	8-21
Enabling IP Preserve	8-22

Configuring a Single Source IP Address	8-25
Overview	8-25
Specifying the Source IP Address	8-25
The Source IP Selection Policy	8-26
Displaying the Source IP Interface Information	8-29
Error Messages	8-32

9 Time Protocols

Contents	9-1
Overview	9-3
TimeP Time Synchronization	9-3
SNTP Time Synchronization	9-3
Selecting a Time Synchronization Protocol or Turning Off Time Protocol Operation	9-4
General Steps for Running a Time Protocol on the Switch:	9-4
Disabling Time Synchronization	9-5
SNTP: Viewing, Selecting, and Configuring	9-5
Menu: Viewing and Configuring SNTP	9-6
CLI: Viewing and Configuring SNTP	9-9
Viewing the Current SNTP Configuration	9-9
Configuring (Enabling or Disabling) the SNTP Mode	9-11
SNTP Client Authentication	9-16
Requirements	9-17
Configuring the Key-Identifier, Authentication Mode, and Key-Value	9-18
Configuring a Trusted Key	9-19
Associating a Key with an SNTP Server	9-20
Enabling SNTP Client Authentication	9-20
Configuring Unicast and Broadcast Mode	9-21
Displaying SNTP Configuration Information	9-21
Saving Configuration Files and the Include-Credentials Command	9-23
TimeP: Viewing, Selecting, and Configuring	9-26
Menu: Viewing and Configuring TimeP	9-27
CLI: Viewing and Configuring TimeP	9-28

Viewing the Current TimeP Configuration	9-29
Configuring (Enabling or Disabling) the TimeP Mode	9-30
SNTP Unicast Time Polling with Multiple SNTP Servers	9-35
Displaying All SNTP Server Addresses Configured on the Switch ..	9-35
Adding and Deleting SNTP Server Addresses	9-36
Menu: Operation with Multiple SNTP Server Addresses Configured	9-36
SNTP Messages in the Event Log	9-36

10 Port Status and Configuration

Contents	10-1
Overview	10-3
Viewing Port Status and Configuring Port Parameters	10-3
Menu: Port Configuration	10-6
CLI: Viewing Port Status and Configuring Port Parameters	10-8
Viewing Port Status and Configuration	10-8
Customizing the Show Interfaces Command	10-10
Error Messages	10-12
Note on Using Pattern Matching with the “Show Interfaces Custom” Command	10-13
Viewing Port Utilization Statistics	10-13
Viewing Transceiver Status	10-14
Enabling or Disabling Ports and Configuring Port Mode	10-15
Enabling or Disabling the USB Port	10-17
Behavior of Autorun When USB Port is Disabled	10-18
Enabling or Disabling Flow Control	10-18
Configuring a Broadcast Limit on the Switch	10-20
Configuring ProCurve Auto-MDIX	10-21
Web: Viewing Port Status and Configuring Port Parameters	10-24
Using Friendly (Optional) Port Names	10-25
Configuring and Operating Rules for Friendly Port Names	10-25
Configuring Friendly Port Names	10-26
Displaying Friendly Port Names with Other Port Data	10-27
Configuring Transceivers and Modules That Haven’t Been Inserted	10-31

Transceivers	10-31
Modules	10-31
Clearing the Module Configuration	10-31
Operating Notes	10-32
Uni-Directional Link Detection (UDLD)	10-33
Configuring UDLD	10-34
Enabling UDLD	10-35
Changing the Keepalive Interval	10-36
Changing the Keepalive Retries	10-36
Configuring UDLD for Tagged Ports	10-36
Viewing UDLD Information	10-37
Configuration Warnings and Event Log Messages	10-39

11 Power Over Ethernet (PoE/PoE+) Operation

Contents	11-1
Introduction to PoE	11-3
PoE Terminology	11-3
PoE Operation	11-5
Configuration Options	11-5
PD Support	11-6
Power Priority Operation	11-7
When Is Power Allocation Prioritized?	11-7
How Is Power Allocation Prioritized?	11-7
Configuring PoE Operation	11-8
Disabling or Re-Enabling PoE Port Operation	11-8
Enabling Support for Pre-Standard Devices	11-8
Configuring the PoE Port Priority Level	11-9
PoE Priority With Two or More Modules	11-10
Controlling PoE Allocation	11-12
Manually Configuring PoE Power Levels	11-13
Configuring PoE Redundancy (Chassis Switches Only)	11-14
Changing the Threshold for Generating a Power Notice	11-15
PoE/PoE+ Allocation Using LLDP Information	11-18
LLDP with PoE	11-18

Displaying the Switch's Global PoE Power Status	11-19
Displaying PoE Status on All Ports	11-21
Displaying the PoE Status on Specific Ports	11-23
Planning and Implementing a PoE Configuration	11-25
Power Requirements	11-25
Assigning PoE Ports to VLANs	11-26
Applying Security Features to PoE Configurations	11-26
Assigning Priority Policies to PoE Traffic	11-26
PoE Event Log Messages	11-28
“Informational” PoE Event-Log Messages	11-28
“Warning” PoE Event-Log Messages	11-29

12 Port Trunking

Contents	12-1
Overview	12-3
Port Trunk Features and Operation	12-5
Trunk Configuration Methods	12-6
Menu: Viewing and Configuring a Static Trunk Group	12-10
CLI: Viewing and Configuring Port Trunk Groups	12-12
Using the CLI To View Port Trunks	12-12
Using the CLI To Configure a Static or Dynamic Trunk Group ...	12-15
Web: Viewing Existing Port Trunk Groups	12-18
Trunk Group Operation Using LACP	12-19
Default Port Operation	12-22
LACP Notes and Restrictions	12-23
Distributed Trunking	12-27
Overview	12-27
Distributed Trunking Interconnect Protocol (DTIP)	12-29
Configuring Distributed Trunking	12-30
ISC Port Configuration	12-30
Distributed Trunking Port Configuration	12-30
Displaying Distributed Trunking Information	12-31
Maximum DT Trunks and Links Supported	12-32

Forwarding Traffic with Distributed Trunking and Spanning Tree	12-32
Forwarding Unicast Traffic Upstream	12-32
Forwarding Broadcast, Multicast, and Unknown Traffic Upstream	12-33
Forwarding Unicast Traffic Downstream (to the Server)	12-33
Forwarding Broadcast, Multicast, and Unknown Traffic Downstream (to the Server)	12-33
Distributed Trunking Restrictions	12-35
Trunk Group Operation Using the “Trunk” Option	12-36
How the Switch Lists Trunk Data	12-37
Outbound Traffic Distribution Across Trunked Links	12-37

13 Port Traffic Controls

Contents	13-1
Overview	13-3
Rate-Limiting	13-4
All Traffic Rate-Limiting	13-4
Configuring Rate-Limiting	13-5
Displaying the Current Rate-Limit Configuration	13-6
Operating Notes for Rate-Limiting	13-8
ICMP Rate-Limiting	13-10
Terminology	13-11
Guidelines for Configuring ICMP Rate-Limiting	13-12
Configuring ICMP Rate-Limiting	13-13
Using Both ICMP Rate-Limiting and All-Traffic Rate-Limiting on the Same Interface	13-14
Displaying the Current ICMP Rate-Limit Configuration	13-14
Operating Notes for ICMP Rate-Limiting	13-15
ICMP Rate-Limiting Trap and Event Log Messages	13-17
Configuring Inbound Rate-Limiting for Broadcast and Multicast Traffic	13-19
Operating Notes	13-21
Guaranteed Minimum Bandwidth (GMB)	13-22
Introduction	13-22

Terminology	13-22
GMB Operation	13-22
Impacts of QoS Queue Configuration on GMB Operation	13-24
Configuring Guaranteed Minimum Bandwidth for Outbound Traffic	13-25
Displaying the Current Guaranteed Minimum Bandwidth Configuration	13-28
GMB Operating Notes	13-29
Jumbo Frames	13-30
Terminology	13-30
Operating Rules	13-31
Configuring Jumbo Frame Operation	13-32
Overview	13-32
Viewing the Current Jumbo Configuration	13-33
Enabling or Disabling Jumbo Traffic on a VLAN	13-35
Configuring a Maximum Frame Size	13-35
Configuring IP MTU	13-36
SNMP Implementation	13-36
Displaying the Maximum Frame Size	13-37
Operating Notes for Maximum Frame Size	13-37
Operating Notes for Jumbo Traffic-Handling	13-37
Troubleshooting	13-40

14 Configuring for Network Management Applications

Contents	14-1
Using SNMP Tools To Manage the Switch	14-3
Overview	14-3
SNMP Management Features	14-4
Configuring for SNMP version 1 and 2c Access to the Switch	14-4
Configuring for SNMP Version 3 Access to the Switch	14-5
SNMP Version 3 Commands	14-6
Enabling SNMPv3	14-7
SNMPv3 Users	14-7
Group Access Levels	14-11
SNMPv3 Communities	14-11

Menu: Viewing and Configuring non-SNMP version 3	
Communities	14-13
CLI: Viewing and Configuring SNMP Community Names	14-15
SNMP Notifications	14-17
Supported Notifications	14-17
General Steps for Configuring SNMP Notifications	14-18
SNMPv1 and SNMPv2c Traps	14-19
Configuring an SNMP Trap Receiver	14-19
Enabling SNMPv2c Informs	14-21
Configuring SNMPv3 Notifications	14-23
Managing Network Security Notifications	14-26
Enabling Link-Change Traps	14-28
Configuring the Source IP Address for SNMP Notifications . .	14-29
Displaying SNMP Notification Configuration	14-31
Configuring Listening Mode	14-33
Advanced Management: RMON	14-34
CLI-Configured sFlow with Multiple Instances	14-34
Terminology	14-34
Configuring sFlow	14-35
Viewing sFlow Configuration and Status	14-35
LLDP (Link-Layer Discovery Protocol)	14-38
Terminology	14-39
General LLDP Operation	14-41
LLDP-MED	14-41
Packet Boundaries in a Network Topology	14-41
Configuration Options	14-42
Options for Reading LLDP Information Collected by the Switch . .	14-44
LLDP and LLDP-MED Standards Compatibility	14-44
LLDP Operating Rules	14-45
Configuring LLDP Operation	14-46
Viewing the Current Configuration	14-46
Configuring Global LLDP Packet Controls	14-48
Configuring SNMP Notification Support	14-52
Configuring Per-Port Transmit and Receive Modes	14-53
Configuring Basic LLDP Per-Port Advertisement Content	14-54

Configuring Support for Port Speed and Duplex Advertisements	14-56
LLDP-MED (Media-Endpoint-Discovery)	14-57
LLDP-MED Topology Change Notification	14-60
LLDP-MED Fast Start Control	14-62
Advertising Device Capability, Network Policy, PoE Status and Location Data	14-62
Configuring Location Data for LLDP-MED Devices	14-66
Displaying Advertisement Data	14-71
Displaying Switch Information Available for Outbound Advertisements	14-72
Displaying LLDP Statistics	14-76
LLDP Operating Notes	14-78
LLDP and CDP Data Management	14-80
LLDP and CDP Neighbor Data	14-80
CDP Operation and Commands	14-82

15 Redundancy (Switches 8200zl)

Contents	15-1
Overview	15-3
Terminology	15-3
How the Management Modules Interact	15-4
Using Redundant Management	15-5
Displaying Redundancy Status	15-5
Enabling or Disabling Redundant Management	15-6
Directing the Standby Module to Become Active	15-8
Setting the Active Management Module for Next Boot	15-9
Enabling and Disabling Fabric Modules	15-12
Management Module Switchover	15-13
Events that Cause a Switchover	15-13
When Switchover Will not Occur	15-13
Consequences of Switchover	15-13
Resetting the Management Module	15-14
Hotswapping Management Modules	15-15
Hotswapping Out the Active Management Module	15-15

When the Standby Module is not Available	15-16
Hotswapping In a Management Module	15-16
Software Version Mismatch Between Active and Hotswapped Module	15-16
Downloading a New Software Version	15-17
File Synchronization after Downloading	15-17
Potential Software Version Mismatches After Downloading	15-18
Downloading a Software Version Serially if the Management Module is Corrupted	15-21
Turning Off Redundant Management	15-21
Disabling Redundancy with Two Modules Present	15-21
Disabling Redundancy With Only One Module Present	15-22
Displaying Management Information	15-23
Active Management Module Commands	15-23
Show Modules	15-23
Show Redundancy	15-24
Show Flash	15-25
Show Version	15-25
Show Log	15-26
Standby Management Module Commands	15-27
Show Redundancy	15-27
Show Flash	15-27
Show Version	15-28
Existing CLI Commands Affected by Redundant Management .	15-29
Boot Command	15-29
Setting the Default Flash for Boot	15-31
Reload Command	15-32
Additional Commands Affected by Redundant Management	15-34
Using the Web Browser for Redundant Management	15-36
Identity Page	15-36
Overview Page	15-37
Redundancy Status Page	15-37
Device View Page	15-38
Management Module LED Behavior	15-40

Active (Actv) LED Behavior	15-40
Standby Led Behavior	15-40
Logging Messages	15-41
Log File	15-41
Crash Files	15-42
Displaying Saved Crash Information	15-42
Notes on How the Active Module is Determined	15-44
Diagram of Decision Process	15-45
Event Log Messages	15-46

A File Transfers

Contents	A-1
Overview	A-4
Downloading Switch Software	A-4
General Software Download Rules	A-5
Using TFTP To Download Software from a Server	A-5
Menu: TFTP Download from a Server to Primary Flash	A-6
CLI: TFTP Download from a Server to Flash	A-8
Enabling TFTP	A-10
Using Auto-TFTP	A-11
Using Secure Copy and SFTP	A-12
How It Works	A-13
The SCP/SFTP Process	A-13
Disable TFTP and Auto-TFTP for Enhanced Security	A-14
Command Options	A-15
Authentication	A-16
SCP/SFTP Operating Notes	A-16
Troubleshooting SSH, SFTP, and SCP Operations	A-18
Using Xmodem to Download Switch Software From a PC or UNIX Workstation	A-20
Menu: Xmodem Download to Primary Flash	A-20
CLI: Xmodem Download from a PC or UNIX Workstation to Primary or Secondary Flash	A-21
Using USB to Transfer Files to and from the Switch	A-22

Using USB to Download Switch Software	A-23
Switch-to-Switch Download	A-24
Menu: Switch-to-Switch Download to Primary Flash	A-25
CLI: Switch-To-Switch Downloads	A-26
Using PCM+ to Update Switch Software	A-27
Copying Software Images	A-28
TFTP: Copying a Software Image to a Remote Host	A-28
Xmodem: Copying a Software Image from the Switch to a Serially Connected PC or UNIX Workstation	A-28
USB: Copying a Software Image to a USB Device	A-29
Transferring Switch Configurations	A-29
TFTP: Copying a Configuration File to a Remote Host	A-30
TFTP: Copying a Configuration File from a Remote Host	A-31
TFTP: Copying a Customized Command File to a Switch	A-31
Xmodem: Copying a Configuration File to a Serially Connected PC or UNIX Workstation	A-33
Xmodem: Copying a Configuration File from a Serially Connected PC or UNIX Workstation	A-33
USB: Copying a Configuration File to a USB Device	A-35
USB: Copying a Configuration File from a USB Device	A-35
Transferring ACL Command Files	A-36
TFTP: Uploading an ACL Command File from a TFTP Server	A-36
Xmodem: Uploading an ACL Command File from a Serially Connected PC or UNIX Workstation	A-38
USB: Uploading an ACL Command File from a USB Device ..	A-38
Copying Diagnostic Data to a Remote Host, USB Device, PC or UNIX Workstation	A-39
Copying Command Output to a Destination Device	A-40
Copying Event Log Output to a Destination Device	A-41
Copying Crash Data Content to a Destination Device	A-41
Copying Crash Log Data Content to a Destination Device	A-43
Enabling or Disabling the USB Port	A-45
Behavior of Autorun When USB Port is Disabled	A-46
Software Versions K.13.XX Operation	A-46
Software Version K.14.XX Operation	A-46

Using USB Autorun	A-47
How It Works	A-47
Security Considerations	A-48
Troubleshooting Autorun Operations	A-49
Configuring Autorun on the Switch	A-50
Enabling Secure Mode	A-50
Operating Notes and Restrictions	A-51
Autorun and Configuring Passwords	A-51
Viewing Autorun Configuration Information	A-52

B Monitoring and Analyzing Switch Operation

Contents	B-1
Overview	B-4
Status and Counters Data	B-5
Menu Access To Status and Counters	B-6
General System Information	B-7
Menu Access	B-7
CLI Access to System Information	B-8
Task Monitor—Collecting Processor Data	B-9
Switch Management Address Information	B-10
Menu Access	B-10
CLI Access	B-11
Module Information	B-12
Menu: Displaying Port Status	B-12
CLI Access	B-13
Port Status	B-14
Menu: Displaying Port Status	B-14
CLI Access	B-15
Web Access	B-15
Viewing Port and Trunk Group Statistics and Flow Control Status	B-15
Menu Access to Port and Trunk Statistics	B-17
CLI Access To Port and Trunk Group Statistics	B-18
Web Browser Access To View Port and Trunk Group Statistics	B-19
Viewing the Switch’s MAC Address Tables	B-19
Menu Access to the MAC Address Views and Searches	B-19

CLI Access for MAC Address Views and Searches	B-22
Spanning Tree Protocol (MSTP) Information	B-23
CLI Access to MSTP Data	B-23
Internet Group Management Protocol (IGMP) Status	B-24
VLAN Information	B-25
Web Browser Interface Status Information	B-27
Traffic Mirroring	B-28
Mirroring Terminology	B-30
Mirrored Traffic Destinations	B-33
Local Destinations	B-33
Remote Destinations	B-33
Monitored Traffic Sources	B-33
Criteria for Selecting Mirrored Traffic	B-34
Mirroring Session Limits	B-34
Mirroring Sessions	B-34
Mirroring Configuration	B-35
Remote Mirroring Endpoint and Intermediate Devices	B-36
Migration to Release K.12.xx	B-37
Migration to Release K.14.01 or Greater	B-37
Using the Menu or Web Interface To Configure Local Mirroring ..	B-39
Menu and Web Interface Limits	B-39
Configuration Steps	B-40
CLI: Configuring Local and Remote Mirroring	B-43
Local Mirroring Overview	B-44
Remote Mirroring Overview	B-46
1. Determine the Mirroring Session and Destination	B-49
2. Configure a Mirroring Destination on a Remote Switch	B-50
3. Configure a Mirroring Session on the Source Switch	B-52
4. Configure the Monitored Traffic in a Mirror Session	B-55
Traffic Selection Options	B-55
Mirroring-Source Restrictions	B-56
Selecting All Inbound/Outbound Traffic to Mirror	B-57
Port Interface with Traffic Direction as the Selection Criteria	B-57
Untagged Mirrored Packets	B-59
VLAN Interface with Traffic Direction as the Selection Criteria	B-61

Selecting Inbound Traffic Using an ACL (Deprecated)	B-62
Selecting Inbound/Outbound Traffic Using a MAC Address	B-63
Selecting Inbound Traffic Using Advanced Classifier-Based Mirroring	B-66
Classifier-Based Mirroring Configuration	B-67
Viewing a Classifier-Based Mirroring Configuration	B-73
Classifier-Based Mirroring Restrictions	B-73
Applying Multiple Mirroring Sessions to an Interface	B-75
Displaying a Mirroring Configuration	B-76
Displaying All Mirroring Sessions Configured on the Switch	B-76
Displaying the Remote Endpoints Configured on the Switch	B-78
Displaying the Mirroring Configuration for a Specific Session	B-79
Displaying Resource Usage for Mirroring Policies	B-84
Viewing the Mirroring Configurations in the Running Configuration File	B-86
Mirroring Configuration Examples	B-87
Example: Local Mirroring Using Traffic-Direction Criteria	B-87
Example: Remote Mirroring Using a Classifier-Based Policy	B-88
Example: Remote Mirroring Using Traffic-Direction Criteria	B-90
Maximum Supported Frame Size	B-92
Enabling Jumbo Frames To Increase the Mirroring Path MTU	B-93
Effect of Downstream VLAN Tagging on Untagged, Mirrored Traffic	B-94
Operating Notes for Traffic Mirroring	B-95
Troubleshooting Traffic Mirroring	B-97

C Troubleshooting

Contents	C-1
Overview	C-4
Troubleshooting Approaches	C-5
Browser or Telnet Access Problems	C-6
Unusual Network Activity	C-8
General Problems	C-8
802.1Q Prioritization Problems	C-9
ACL Problems	C-9

IGMP-Related Problems	C-14
LACP-Related Problems	C-14
Mesh-Related Problems	C-15
Port-Based Access Control (802.1X)-Related Problems	C-15
QoS-Related Problems	C-18
Radius-Related Problems	C-18
Spanning-Tree Protocol (MSTP) and Fast-Uplink Problems	C-19
SSH-Related Problems	C-20
TACACS-Related Problems	C-22
TimeP, SNTP, or Gateway Problems	C-24
VLAN-Related Problems	C-24
Fan Failure	C-26
Using the Event Log for Troubleshooting Switch Problems	C-27
Event Log Entries	C-27
Menu: Displaying and Navigating in the Event Log	C-35
CLI: Displaying the Event Log	C-36
CLI: Clearing Event Log Entries	C-36
CLI: Turning Event Numbering On	C-37
Using Log Throttling to Reduce Duplicate Event Log and SNMP Messages	C-37
Log Throttle Periods	C-38
Example of Log Throttling	C-38
Example of Event Counter Operation	C-40
Debug/Syslog Operation	C-41
Debug/Syslog Messaging	C-41
Debug/Syslog Destination Devices	C-41
Debug/Syslog Configuration Commands	C-42
Configuring Debug/Syslog Operation	C-44
Displaying a Debug/Syslog Configuration	C-46
Debug Command	C-50
Debug Messages	C-50
Debug Destinations	C-52
Logging Command	C-54
Configuring a Syslog Server	C-55
Adding a Description for a Syslog Server	C-57

Adding a Priority Description	C-58
Configuring the Severity Level for Event Log Messages Sent to a Syslog Server	C-59
Configuring the System Module Used to Select the Event Log Messages Sent to a Syslog Server	C-60
Operating Notes for Debug and Syslog	C-60
Diagnostic Tools	C-62
Port Auto-Negotiation	C-63
Ping and Link Tests	C-63
Web: Executing Ping or Link Tests	C-64
CLI: Ping Test	C-65
Link Tests	C-66
Traceroute Command	C-67
Viewing Switch Configuration and Operation	C-71
CLI: Viewing the Startup or Running Configuration File	C-71
Web: Viewing the Configuration File	C-71
CLI: Viewing a Summary of Switch Operational Data	C-72
Saving show tech Command Output to a Text File	C-73
Customizing show tech Command Output	C-75
CLI: Viewing More Information on Switch Operation	C-78
Pattern Matching When Using the Show Command	C-79
CLI: Useful Commands for Troubleshooting Sessions	C-82
Restoring the Factory-Default Configuration	C-83
CLI: Resetting to the Factory-Default Configuration	C-83
Clear/Reset: Resetting to the Factory-Default Configuration	C-83
Restoring a Flash Image	C-84
DNS Resolver	C-87
Terminology	C-87
Basic Operation	C-88
Configuring and Using DNS Resolution with DNS-Compatible Commands	C-89
Configuring a DNS Entry	C-90
Example Using DNS Names with Ping and Traceroute	C-91
Viewing the Current DNS Configuration	C-93
Operating Notes	C-94

Event Log Messages	C-95
Locator LED (Locating a Switch)	C-96

D MAC Address Management

Contents	D-1
Overview	D-2
Determining MAC Addresses	D-3
Menu: Viewing the Switch's MAC Addresses	D-4
CLI: Viewing the Port and VLAN MAC Addresses	D-5
Viewing the MAC Addresses of Connected Devices	D-7

E Monitoring Resources

Contents	E-1
Viewing Information on Resource Usage	E-2
Policy Enforcement Engine	E-2
Displaying Current Resource Usage	E-4
When Insufficient Resources Are Available	E-7

F Daylight Savings Time on ProCurve Switches

G Scalability: IP Address, VLAN, and Routing Maximum Values

H Switch Licensing

General Procedure	H-1
-------------------------	-----

I Power-Saving Features

Contents	I-1
Overview	I-2
Configuring the Power-Saving Options	I-3
Configuring the Savepower module Option	I-3
Configuring the Savepower LED Option	I-4
Configuring the Savepower port-low-pwr Option	I-6

Show Savepower Commands I-6

J Network Out-of-Band Management (OOBM) for the 6600 Switch

Contents J-1

Concepts J-2

 Example J-4

 OOBM and Switch Applications J-5

Tasks J-6

 OOBM Configuration J-6

 OOBM Context J-6

 OOBM Enable/disable J-7

 OOBM Port Enable/disable J-8

 OOBM Port Speed Control J-9

 OOBM IPv4 Address Configuration J-10

 OOBM IPv4 Default Gateway Configuration J-10

 OOBM Show Commands J-11

 Show OOBM J-11

 Show OOBM IP Configuration J-12

 Show OOBM ARP Information J-12

 Application Server Commands J-13

 Application Client Commands J-14

 Example J-15

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/manuals.

Printed Publications

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

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

Electronic Publications

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

- *Management and Configuration Guide*—Describes how to configure, manage, and monitor basic switch operation.
- *Advanced Traffic Management Guide*—Explains how to configure traffic management features such as 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.
- *IPv6 Configuration Guide*—Describes the IPv6 protocol operations that are supported on the switch.
- *Command Line Interface Reference Guide*—Provides a comprehensive description of CLI commands, syntax, and operations.
- *Event Log Message Reference Guide*—Provides a comprehensive description of event log messages.
- *Release Notes*—Describe new features, fixes, and enhancements that become available between revisions of the main product guide.

Software Feature Index

For the software manual set supporting your 3500/3500yl/5400zl/6200yl/6600/8200zl switch model, this feature index indicates which manual to consult for information on a given software feature.

Note

This Index does not cover IPv6 capable software features. For information on IPv6 protocol operations and features (such as DHCPv6, DNS for IPv6, Ping6, and MLD Snooping), refer to the *IPv6 Configuration Guide*.

Intelligent Edge Software Features. These features are automatically included on all switches.

Premium License Software Features. For the HP ProCurve 3500, 3500yl, 5400zl, 6600, and 8200zl switches, Premium License features can be acquired by purchasing the optional Premium License and installing it on the Intelligent Edge version of these switches. (These features are automatically included on the HP ProCurve 6200yl switches.)

Premium License 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	
QinQ (Provider Bridging)		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			

Intelligent Edge Software Features	Manual			
	Management and Configuration	Advanced Traffic Management	Multicast and Routing	Access Security Guide
802.1X Multiple Authenticated Clients Per Port				X
Access Control Lists (ACLs)				X
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			
Core Dump	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			
Distributed Trunking	X			
Downloading Software	X			
Dynamic ARP Protection				X
Dynamic Configuration Arbiter				X
Dynamic IP Lockdown				X
Eavesdrop Protection				X
Equal Cost Multi-Path (ECMP)			X	
Event Log	X			

Intelligent Edge Software Features	Manual			
	Management and Configuration	Advanced Traffic Management	Multicast and Routing	Access Security Guide
Factory Default Settings	X			
Flow Control (802.3x)	X			
File Management	X			
File Transfers	X			
Friendly Port Names	X			
Guaranteed Minimum Bandwidth (GMB)	X			
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			
Key Management System (KMS)				X
LACP	X			
LLDP	X			
LLDP-MED	X			
Loop Protection		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

Intelligent Edge Software Features	Manual			
	Management and Configuration	Advanced Traffic Management	Multicast and Routing	Access Security Guide
Multiple Configuration Files	X			
Network Management Applications (SNMP)	X			
Out-of-Band Management (OOBM)	X			
OpenView Device Management	X			
Passwords and Password Clear Protection				X
ProCurve Manager (PCM)	X			
Ping	X			
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 and 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	
SavePower Features	X			
Secure Copy	X			

Intelligent Edge Software Features	Manual			
	Management and Configuration	Advanced Traffic Management	Multicast and Routing	Access Security Guide
sFlow	X			
SFTP	X			
SNMPv3	X			
Software Downloads (SCP/SFTP, TFTP, Xmodem)	X			
Source-Port Filters				X
Spanning Tree (STP, RSTP, MSTP)		X		
SSHv2 (Secure Shell) Encryption				X
SSL (Secure Socket Layer)				X
Stacking (3500/3500yl/6200yl/6600 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

Intelligent Edge Software Features	Manual			
	Management and Configuration	Advanced Traffic Management	Multicast and Routing	Access Security Guide
Web-based Authentication				X
Web UI	X			

Getting Started

Contents

Introduction	1-2
Conventions	1-2
Command Syntax Statements	1-2
Command Prompts	1-3
Screen Simulations	1-3
Configuration and Operation Examples	1-3
Keys	1-3
Sources for More Information	1-4
Getting Documentation From the Web	1-6
Online Help	1-6
Menu Interface	1-6
Command Line Interface	1-7
Web Browser Interface	1-7
Need Only a Quick Start?	1-8
IP Addressing	1-8
To Set Up and Install the Switch in Your Network	1-8
Physical Installation	1-8

Introduction

This guide is intended for use with the following ProCurve switches:

- 8200zl switches
- 6600 switches
- 5400zl switches
- 3500, 3500yl and 6200yl switches

It describes how to use the command line interface (CLI), Menu interface, and web browser to configure, manage, monitor, and troubleshoot switch operation. For an overview of product documentation for the above switches, refer to “*Product Documentation*” on page xiii. To download the switch documentation, visit the ProCurve Networking manuals web page at www.hp.com/go/procurve/manuals.

Conventions

This guide uses the following conventions for commands and screen displays.

Command Syntax Statements

Syntax: ip < default-gateway < *ip-addr* >> | routing >

Syntax: show interfaces [*port-list*]

- Vertical bars (|) separate alternative, mutually exclusive elements.
- Square brackets ([]) indicate optional elements.
- Braces (< >) enclose required elements.
- Braces within square brackets ([< >]) indicate a required element within an optional choice.
- Boldface indicates use of a CLI command, part of a CLI command syntax, or other displayed element in general text. For example:
 “Use the **copy tftp** command to download the key from a TFTP server.”
- Italics indicate variables for which you must supply a value when executing the command. For example, in this command syntax, you must provide one or more port numbers:

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

Command Prompts

In the default configuration, your switch displays a CLI prompt similar to the following example:

```
ProCurve 8212z1#
```

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

```
ProCurve#
```

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

Screen Simulations

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

```
ProCurve> show version
Image stamp:   /sw/code/build/info
               March 1, 2007 13:43:13
               K.12.01
               139
ProCurve>
```

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

In some cases, brief command-output sequences appear without figure identification. For example:

```
ProCurve(config)# clear public-key
ProCurve(config)# show ip client-public-key
show_client_public_key: cannot stat keyfile
```

Configuration and Operation Examples

Unless otherwise noted, examples using a particular switch model apply to all switch models covered by this guide.

Keys

Simulations of actual keys use a bold, sans-serif typeface with square brackets. For example, the Tab key appears as **[Tab]** and the “Y” key appears as **[Y]**.

Sources for More Information

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

- **Feature Index**—For information on which manual to consult for a given software feature, refer to the “Software Feature Index” on page xiv.

Note

For the latest version of all HP ProCurve switch documentation referred to below, including Release Notes covering recently added features, visit the ProCurve Networking manuals web page at www.hp.com/go/procurve/manuals.

- **Software Release Notes**—*Release Notes* are posted on the HP ProCurve Networking web site and provide information on new software updates:
 - new features and how to configure and use them
 - software management, including downloading software to the switch
 - software fixes addressed in current and previous releases
- **Product Notes and Software Update Information**—The printed *Read Me First* shipped with your switch provides software update information, product notes, and other information.
- **Installation and Getting Started Guide**—Use the *Installation and Getting Started Guide* shipped with your switch to prepare for and perform the physical installation. This guide also steps you through connecting the switch to your network and assigning IP addressing, as well as describing the LED indications for correct operation and trouble analysis.
- **Management and Configuration Guide**—Use this guide for information on topics such as:
 - various interfaces available on the switch
 - memory and configuration operation
 - interface access
 - IP addressing
 - time protocols
 - port configuration, trunking, traffic control, and PoE operation
 - Redundant management
 - SNMP, LLDP, and other network management topics

- file transfers, switch monitoring, troubleshooting, and MAC address management
- *Advanced Traffic Management Guide*—Use this guide for information on topics such as:
 - VLANs: Static port-based and protocol VLANs, and dynamic GVRP VLANs
 - spanning-Tree: 802.1D (STP), 802.1w (RSTP), and 802.1s (MSTP)
 - meshing
 - Quality-of-Service (QoS)
 - Access Control Lists (ACLs)
 - Out-of-Band Management (6600)
- *Multicast and Routing Guide*—Use this guide for information on topics such as:
 - IGMP
 - PIM (SM and DM)
 - IP routing
 - VRRP
- *Access Security Guide*—Use this guide for information on topics such as:
 - Local username and password security
 - Web-Based and MAC-based authentication
 - RADIUS and TACACS+ authentication
 - SSH (Secure Shell) and SSL (Secure Socket Layer) operation
 - 802.1X access control
 - Port security operation with MAC-based control
 - Authorized IP Manager security
 - Key Management System (KMS)
- *IPv6 Configuration Guide*—Use this guide for information on topics such as:
 - Overview of IPv6 operation and features supported in software release K.13.01 or greater
 - Configuring IPv6 addressing
 - Using IPv6 management, security, and troubleshooting features

Getting Documentation From the Web

To obtain the latest versions of documentation and release notes for your switch, go to the ProCurve Networking manuals web page at www.hp.com/go/procurve/manuals.

Online Help

Menu Interface

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

```
----- CONSOLE - MANAGER MODE -----  
Switch Configuration - Internet (IP) Service  
  
IP Routing : Disabled  
  
Default Gateway :  
Default TTL    : 64  
Arp Age       : 20  
  
IP Config [DHCP/Bootp] : Manual  
IP Address    : 10.35.204.104  
Subnet Mask   : 255.255.240.0  
  
Actions->  Cancel    Edit    Save    Help  
  
Display help information.  
Use arrow keys to change action selection and <Enter> to execute action.
```

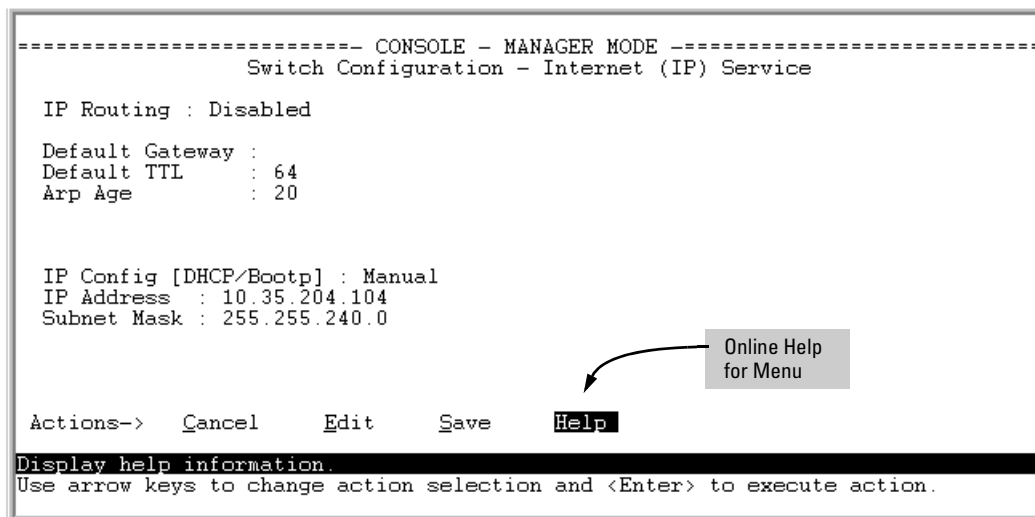


Figure 1-2. Online Help for Menu Interface

Command Line Interface

If you need information on a specific command in the CLI, type the command name followed by **help**. For example:

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

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

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

Figure 1-3. Example of CLI Help

Web Browser Interface

If you need information on specific features in the HP ProCurve Web Browser Interface (hereafter referred to as the “web browser interface”), use the online Help. You can access the Help by clicking on the “Help” text in the lower right corner of any of the web browser interface screens.




Figure 1-4. Web Browser Interface Online Help

Note

To access the online Help for the Web browser interface, you need either ProCurve Manager (version 1.5 or greater) installed on your network or an active connection to the World Wide Web. Otherwise, Online help for the web browser interface will not be available.

Need Only a Quick Start?

IP Addressing

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

- Enter **setup** at the CLI Manager level prompt.

```
Procurve# setup
```

- In the Main Menu of the Menu interface, select

8. Run Setup

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

To Set Up and Install the Switch in Your Network

Physical Installation

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

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

For the latest version of the *Installation and Getting Started Guide* for your switch, refer to “Getting Documentation From the Web” on page 1-6.

Selecting a Management Interface

Contents

Overview	2-2
Understanding Management Interfaces	2-2
Advantages of Using the Menu Interface	2-3
Advantages of Using the CLI	2-4
General Benefits	2-4
Information on Using the CLI	2-4
Advantages of Using the Web Browser Interface	2-5
Advantages of Using ProCurve Manager or ProCurve Manager Plus	2-7
Custom Login Banners for the Console and Web Browser Interfaces	2-9
Banner Operation with Telnet, Serial, or SSHv2 Access	2-9
Banner Operation with Web Browser Access	2-9
Configuring and Displaying a Non-Default Banner	2-10
Example of Configuring and Displaying a Banner	2-11
Operating Notes	2-13

Overview

This chapter describes the following:

- Management interfaces for the switches covered in this guide
 - Advantages of using each interface
-

Understanding Management Interfaces

Management interfaces enable you to reconfigure the switch and to monitor switch status and performance. The switch offers the following interfaces:

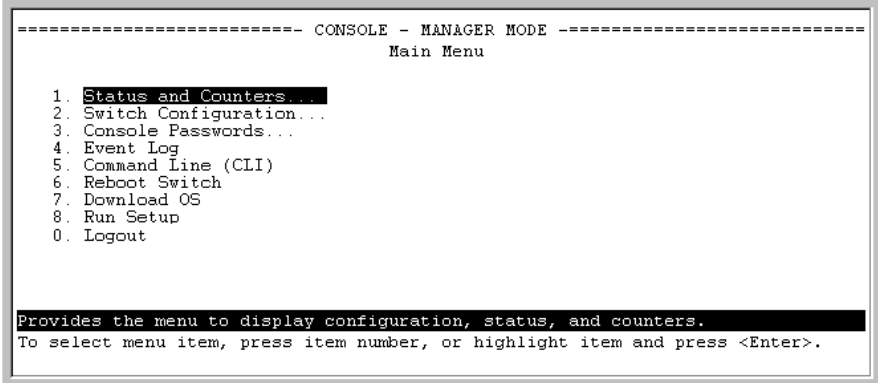
- **Menu interface**—a menu-driven interface offering a subset of switch commands through the built-in VT-100/ANSI console—**2-3**
- **CLI**—a command line interface offering the full set of switch commands through the VT-100/ANSI console built into the switch—**2-4**
- **Web browser interface**—a switch interface offering status information and a subset of switch commands through a standard web browser (such as Netscape Navigator or Microsoft Internet Explorer)—**2-5**
- **ProCurve Manager (PCM)**—a windows-based network management solution included in-box with all manageable ProCurve devices. Features include automatic device discovery, network status summary, topology and mapping, and device management.
- **ProCurve Manager Plus (PCM+)**—a complete windows-based network management solution that provides both the basic features offered with PCM, as well as more advanced management features, including in-depth traffic analysis, group and policy management, configuration management, device software updates, and advanced VLAN management. (ProCurve includes a copy of PCM+ in-box for a free 30-day trial.)

This manual describes how to use the menu interface (Chapter 3), the CLI (Chapter 4), the web browser interface (Chapter), and how to use these interfaces to configure and monitor the switch.

For information on how to access the web browser interface Help, see “Online Help for the Web Browser Interface” on page 5-12.

To use ProCurve Manager or ProCurve Manager Plus, refer to the *Getting Started Guide* and the *Administrator's Guide*, which are available electronically with the software for these applications. For more information, visit the ProCurve Networking web site at www.procurve.com.

Advantages of Using the Menu Interface



```
----- CONSOLE - MANAGER MODE -----
                          Main Menu

1. Status and Counters...
2. Switch Configuration...
3. Console Passwords...
4. Event Log
5. Command Line (CLI)
6. Reboot Switch
7. Download OS
8. Run Setup
0. Logout

Provides the menu to display configuration, status, and counters.
To select menu item, press item number, or highlight item and press <Enter>.
```

Figure 2-1. Example of the Console Interface Display

- **Provides quick, easy management access** to a menu-driven subset of switch configuration and performance features:

- IP addressing
- VLANs and GVRP
- Port Security
- Port and Static Trunk Group
- Spanning Tree
- System information
- Local passwords
- SNMP communities
- Time protocols

The menu interface also provides access for:

- Setup screen
- Event Log display
- Switch and port status displays
- Switch and port statistic and counter displays
- Reboots
- Software downloads

- **Offers out-of-band access** (through the RS-232 connection) to the switch, so network bottlenecks, crashes, lack of configured or correct IP address, and network downtime do not slow or prevent access

- **Enables Telnet (in-band) access** to the menu functionality.
- **Allows faster navigation**, avoiding delays that occur with slower display of graphical objects over a web browser interface.
- **Provides more security**; configuration information and passwords are not seen on the network.

Advantages of Using the CLI

ProCurve>	Prompt for Operator Level
ProCurve#	Prompt for Manager Level
ProCurve (config) #	Prompt for Global Configuration Level
ProCurve (<context>) #	Prompt for Context Configuration Levels
For example:	
ProCurve (eth-1-5) #	
ProCurve (vlan-1) #	
ProCurve (pim) #	
ProCurve (rip) #	

Figure 2-2. Command Prompt Examples

General Benefits

- Provides access to the complete set of the switch configuration, performance, and diagnostic features.
- Offers out-of-band access (through the RS-232 connection) or Telnet (in-band) access.
- Enables quick, detailed system configuration and management access to system operators and administrators experienced in command prompt interfaces.
- Provides help at each level for determining available options and variables.

Information on Using the CLI

- For information on how to use the CLI, refer to Chapter 4. “Using the Command Line Interface (CLI)”.

- To perform specific procedures (such as configuring IP addressing or VLANs), use the Contents listing at the front of the manual to locate the information you need.
- For monitoring and analyzing switch operation, refer to Appendix B.
- For information on individual CLI commands, refer to the Index or to the online Help provided in the CLI interface.

Advantages of Using the Web Browser Interface

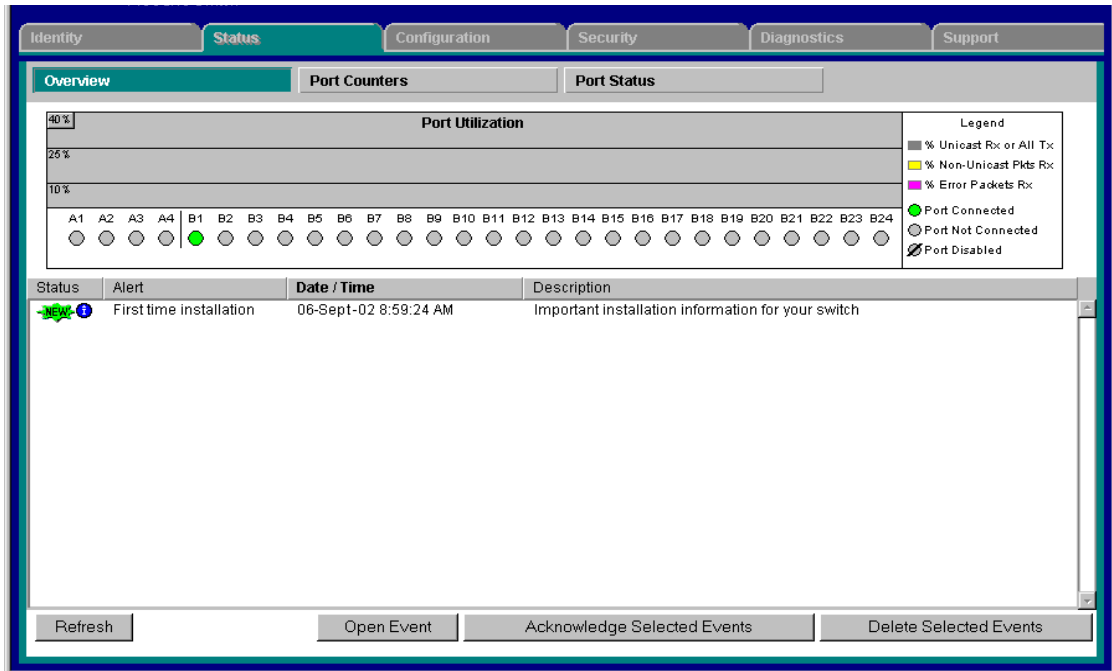


Figure 2-3. Example of the Web Browser Interface

- **Easy access** to the switch from anywhere on the network
- **Familiar browser interface**—locations of window objects consistent with commonly used browsers, uses mouse clicking for navigation, no terminal setup
- **Many features have all their fields in one screen** so you can view all values at once

Selecting a Management Interface

Advantages of Using the Web Browser Interface

- **More visual cues**, using colors, status bars, device icons, and other graphical objects instead of relying solely on alphanumeric values
- **Display of acceptable ranges of values available** in configuration list boxes

Advantages of Using ProCurve Manager or ProCurve Manager Plus

You can operate ProCurve Manager and ProCurve Manager Plus (PCM and PCM+) from a PC on the network to monitor traffic, manage your hubs and switches, and proactively recommend network changes to increase network uptime and optimize performance. Easy to install and use, PCM and PCM+ are the answers to your management challenges.




Figure 2-4. Example of the Home Page for ProCurve Manager Plus

Selecting a Management Interface

Advantages of Using ProCurve Manager or ProCurve Manager Plus

PCM and PCM+ enable greater control, uptime, and performance in your network:

- Features and benefits of ProCurve Manager:
 - **Network Status Summary:** Upon boot-up, a network status screen displays high-level information on network devices, end nodes, events, and traffic levels. From here, users can research any one of these areas to get more details.
 - **Alerts and Troubleshooting:** An events summary screen displays alerts to the user and categorizes them by severity, making it easier to track where bottlenecks and issues exist in the network. Alerts present detailed information on the problem, even down to the specific port.
 - **Automatic Device Discovery:** This feature is customized for fast discovery of all ProCurve manageable network devices. The user can define which IP subnets to discover.
 - **Topology and Mapping:** This feature automatically creates a map of discovered network devices. Maps are color-coded to reflect device status and can be viewed at multiple levels (physical view, subnet view, or VLAN view).
 - **Device Management:** Many device-focused tasks can be performed directly by the software, or the user can access web-browser and command-line interfaces with the click of a button to manage individual devices from inside the tool.
- Features and benefits of ProCurve Manager Plus:
 - **All of the Features of ProCurve Manager:** Refer to the above listing.
 - **In-Depth Traffic Analysis:** An integrated, low-overhead traffic monitor interface shows detailed information on traffic throughout the network. Using enhanced traffic analysis protocols such as Extended RMON and sFlow, users can monitor overall traffic levels, segments with the highest traffic, or even the top users within a network segment.
 - **Group and Policy Management:** Changes in configuration are tracked and logged, and archived configurations can be applied to one or many devices. Configurations can be compared over time or between two devices, with the differences highlighted for users.
 - **Advanced VLAN Management:** A new, easy-to-use VLAN management interface allows users to create and assign VLANs across the entire network, without having to access each network device individually.

- **Device Software Updates:** This feature automatically obtains new device software images from ProCurve and updates devices, allowing users to download the latest version or choose the desired version. Updates can be scheduled easily across large groups of devices, all at user-specified times.
- **Investment Protection:** The modular software architecture of ProCurve Manager Plus will allow ProCurve to offer network administrators add-on software solutions that complement their needs.

Custom Login Banners for the Console and Web Browser Interfaces

You can now configure the switch to display a login banner of up to 3070 characters when an operator initiates a management session with the switch through any of the following methods:

- Telnet
- serial connection
- SSHv2
- Web browser

The default banner displays product registration information; the copyright splash is no longer displayed.

If a banner is configured, the banner page is displayed when you access the Web user interface. The default product registration information is not displayed as there is already a product registration prompt displayed in the Web user interface.

Banner Operation with Telnet, Serial, or SSHv2 Access

When a system operator begins a login session, the switch displays the banner above the local password prompt or, if no password is configured, above the **Press any key to continue prompt**. Entering a correct password or, if no password is configured, pressing any key clears the banner from the CLI and displays the CLI prompt. (Refer to Figure 2-5 on page 2-11.)

Banner Operation with Web Browser Access

When a system operator uses a Web browser to access the switch, the text of a non-default banner configured on the switch appears in a dedicated banner window with a link to the Web agent home page. Clicking on **To Home Page**

clears the banner window and prompts the user for a password (if configured). Following entry of the correct username/password information (or if no username/password is required), the switch then displays either the Registration page or the switch's home page. Note that if the banner feature is disabled or if the switch is using the factory-default banner shown in figure 2-5, then the banner page does not appear in the Web browser when an operator initiates a login session with the switch.

Configuring and Displaying a Non-Default Banner

You can enable or disable banner operation using either the switch's CLI or an SNMP application. The steps include:

1. Enable non-default banner operation and define the endpoint delimiter for the banner.
2. Enter the desired banner text, including any specific line breaks you want.
3. Enter the endpoint delimiter.

Use **show banner motd** to display the current banner status.

Syntax: banner motd < delimiter >
no banner motd

*This command defines the single character used to terminate the banner text and enables banner text input. You can use any character except a blank space as a delimiter. The **no** form of the command disables the login banner feature.*

< banner-text-string >

*The switch allows up to 3070 banner characters, including blank spaces and CR-LF ([Enter]). (The tilde "~" and the delimiter defined by **banner motd <delimiter>** are not allowed as part of the banner text.) While entering banner text, you can backspace to edit the current line (that is, a line that has not been terminated by a CR-LF.) However, terminating a line in a banner by entering a CR-LF prevents any further editing of that line. To edit a line in a banner entry after terminating the line with a CR-LF requires entering the delimiter described above and then re-configuring new banner text.*

The banner text string must terminate with the character defined by banner motd < delimiter >.

Note: In redundant management, the banner is not seen on the standby module, only the active module.

Example of Configuring and Displaying a Banner

Suppose a system operator wanted to configure the following banner message on her company's switches:

```
This is a private system maintained by the
      Allied Widget Corporation.
Unauthorized use of this system can result in
      civil and criminal penalties!
```

In this case, the operator will use the [Enter] key to create line breaks, blank spaces for line centering, and the % symbol to terminate the banner message.

```
ProCurve(config)# banner motd %
Enter TEXT message. End with the character '%'
      This is a private system maintained by the
      Allied Widget Corporation.
      Unauthorized use of this system can result in
      civil and criminal penalties!%
ProCurve(config)# write memory
```

Figure 2-5. Example of Configuring a Login Banner

To view the current banner configuration, use either the **show banner motd** or **show running** command.

```
ProCurve(config)# show banner motd

Banner Information

Banner status: Enabled
Configured Banner:

      This is a private system maintained by the
      Allied Widget Corporation.
      Unauthorized use of this system can result in
      civil and criminal penalties!
```

Figure 2-6. Example of show banner motd Output

Selecting a Management Interface

Advantages of Using ProCurve Manager or ProCurve Manager Plus

```
ProCurve(config)# show running

Running configuration:

; J8697A Configuration Editor; Created on release #K.11.00

hostname "ProCurve"
module 1 type J8702A
module 2 type J8702A
snmp-server community "notpublic" Unrestricted
vlan 1
  name "DEFAULT_VLAN"
  untagged A1-A24,B1-B24
  ip address dhcp-bootp
  _exit_
banner motd "  This is a private system maintained by the
              Allied Widget Corporation.
              Unauthorized use of this system can result in
              civil and criminal penalties!"
password manager
password operator
```

Shows the current banner configuration.

Figure 2-7. The Current Banner Appears in the Switch's Running-Config File

The next time someone logs onto the switch's management CLI, the following appears:

```
Copyright (C) 1991-2005 Hewlett-Packard Co. All Rights Reserved.

                RESTRICTED RIGHTS LEGEND

Use, duplication, or disclosure by the Government is subject to restrictions
as set forth in subdivision (b) (3) (ii) of the Rights in Technical Data and
Computer Software clause at 52.227-7013.

                HEWLETT-PACKARD COMPANY, 3000 Hanover St., Palo Alto, CA 94303

  This is a private system maintained by the
  Allied Widget Corporation.
  Unauthorized use of this system can result in
  civil and criminal penalties!

Password: █
```

The login screen displays the configured banner.
Entering a correct password clears the banner and displays the CLI prompt.

Figure 2-8. Example of CLI Result of the Login Banner Configuration

If someone uses a Web browser to log in to the switch interface, the following message appears:




Figure 2-9. Example of Web Browser Interface Result of the Login Banner Configuration

Operating Notes

- The default banner appears only when the switch is in the factory default configuration. Using **no banner motd** deletes the currently configured banner text and blocks display of the default banner. The default banner is restored only if the switch is reset to its factory-default configuration.
- The switch supports one banner at any time. Configuring a new banner replaces any former banner configured on the switch.
- If the switch is configured with **ssh version 1** or **ssh version 1-or-2**, configuring the banner sets the SSH configuration to ssh version 2 and displays the following message in the CLI:

```
Warning: SSH version has been set to v2.
```

- If a banner is configured, the switch does not allow configuration with **ssh version 1** or **ssh version 1-or-2**. Attempting to do so produces the following error message in the CLI:

```
Banner has to be disabled first.
```

- If a banner is enabled on the switch, the Web browser interface displays the following link to the banner page:

Notice to all users

Selecting a Management Interface

Advantages of Using ProCurve Manager or ProCurve Manager Plus

Using the Menu Interface

Contents

Overview	3-2
Starting and Ending a Menu Session	3-3
How To Start a Menu Interface Session	3-4
How To End a Menu Session and Exit from the Console:	3-5
Main Menu Features	3-7
Screen Structure and Navigation	3-9
Rebooting the Switch	3-12
Menu Features List	3-14
Where To Go From Here	3-15

Overview

This chapter describes the following features:

- Overview of the Menu Interface (page 3-2)
- Starting and ending a Menu session (page 3-3)
- The Main Menu (page 3-7)
- Screen structure and navigation (page 3-9)
- Rebooting the switch (page 3-12)

The menu interface operates through the switch console to provide you with a subset of switch commands in an easy-to-use menu format enabling you to:

- Perform a “quick configuration” of basic parameters, such as the IP addressing needed to provide management access through your network
- Configure these features:
 - Manager and Operator passwords
 - System parameters
 - IP addressing
 - Time protocol
 - Ports
 - Trunk groups
 - A network monitoring port
 - SNMP community names
 - IP authorized managers
 - VLANs (Virtual LANs) and GVRP
- View status, counters, and Event Log information
- Update switch software
- Reboot the switch

For a detailed list of menu features, see the “Menu Features List” on page 3-14.

Privilege Levels and Password Security. ProCurve strongly recommends that you configure a Manager password to help prevent unauthorized access to your network. A Manager password grants full read-write access to the switch. An Operator password, if configured, grants access to status and counter, Event Log, and the Operator level in the CLI. After you configure passwords on the switch and log off of the interface, access to the menu interface (and the CLI and web browser interface) will require entry of either the Manager or Operator password. (If the switch has only a Manager password, then someone without a password can still gain read-only access.)

Note

If the switch has neither a Manager nor an Operator password, anyone having access to the console interface can operate the console with full manager privileges. Also, if you configure only an Operator password, entering the Operator password enables full manager privileges.

For more information on passwords, refer to the *Access Security Guide* for your switch.

Menu Interaction with Other Interfaces.

- The menu interface displays the current running-config parameter settings. You can use the menu interface to save configuration changes made in the CLI only if the CLI changes are in the running config when you save changes made in the menu interface. (For more on how switch memory manages configuration changes, see Chapter 6, “Switch Memory and Configuration”.)
- A configuration change made through any switch interface overwrites earlier changes made through any other interface.
- The Menu Interface and the CLI (Command Line Interface) both use the switch console. To enter the menu from the CLI, use the **menu** command. To enter the CLI from the Menu interface, select **Command Line (CLI)** option.)

Starting and Ending a Menu Session

You can access the menu interface using any of the following:

- A direct serial connection to the switch’s console port, as described in the installation guide you received with the switch
- A Telnet connection to the switch console from a networked PC or the switch’s web browser interface. Telnet requires that an IP address and subnet mask compatible with your network have already been configured on the switch.

Note

This section assumes that either a terminal device is already configured and connected to the switch (see the *Installation and Getting Started Guide* shipped with your switch) or that you have already configured an IP address on the switch (required for Telnet access).

How To Start a Menu Interface Session

In its factory default configuration, the switch console starts with the CLI prompt. To use the menu interface with Manager privileges, go to the Manager level prompt and enter the **menu** command.

1. Use one of these methods to connect to the switch:
 - A PC terminal emulator or terminal
 - Telnet
2. Do one of the following:
 - If you are using Telnet, go to step 3.
 - If you are using a PC terminal emulator or a terminal, press **[Enter]** one or more times until a prompt appears.
3. When the switch screen appears, do one of the following:
 - If a password has been configured, the password prompt appears.

```
Password: _
```

Type the Manager password and press **[Enter]**. Entering the Manager password gives you manager-level access to the switch. (Entering the Operator password gives you operator-level access to the switch. Refer to the *Access Security Guide* for your switch.)
 - If no password has been configured, the CLI prompt appears. Go to the next step.
4. When the CLI prompt appears, display the Menu interface by entering the **menu** command. For example:

```
ProCurve# menu [Enter]
```

results in the following display:

```
----- CONSOLE - MANAGER MODE -----  
Main Menu  
  
1. Status and Counters...  
2. Switch Configuration...  
3. Console Passwords...  
4. Event Log  
5. Command Line (CLI)  
6. Reboot Switch  
7. Download OS  
8. Run Setup  
0. Logout  
  
Provides the menu to display configuration, status, and counters.  
To select menu item, press item number, or highlight item and press <Enter>.
```

Figure 3-1. Example of the Main Menu with Manager Privileges

For a description of Main Menu features, see “Main Menu Features” on page 3-7.

Note

To configure the switch to start with the menu interface instead of the CLI, go to the Manager level prompt in the CLI, enter the **setup** command, and in the resulting display, change the **Logon Default** parameter to **Menu**. For more information, see the *Installation and Getting Started Guide* you received with the switch.

How To End a Menu Session and Exit from the Console:

The method for ending a menu session and exiting from the console depends on whether, during the session, you made any changes to the switch configuration that require a switch reboot to activate. (Most changes via the menu interface need only a **Save**, and do not require a switch reboot.) Configuration changes needing a reboot are marked with an asterisk (*) next to the configured item in the menu and also next to the **Switch Configuration** item in the Main Menu.

Asterisk indicates a configuration change that requires a reboot to activate.

```
----- CONSOLE - MANAGER MODE -----  
Main Menu  
  
1. Status and Counters...  
*2. Switch Configuration...  
3. Console Passwords...  
4. Event Log  
5. Command Line (CLI)  
6. Reboot Switch  
7. Download OS  
8. Run Setup  
0. Logout  
  
Displays the menu for customizing the switch configuration.  
To select menu item, press item number, or highlight item and press <Enter>.  
(*Needs reboot to activate changes.)
```

Figure 3-2. Example Indication of a Configuration Change Requiring a Reboot

1. In the current session, if you have not made configuration changes that require a switch reboot to activate, return to the Main Menu and press [0] (zero) to log out. Then just exit from the terminal program, turn off the terminal, or quit the Telnet session.
2. If you *have* made configuration changes that require a switch reboot—that is, if an asterisk (*) appears next to a configured item or next to **Switch Configuration** in the Main Menu:
 - a. Return to the Main Menu.
 - b. Press [6] to select **Reboot Switch** and follow the instructions on the reboot screen.

Rebooting the switch terminates the menu session, and, if you are using Telnet, disconnects the Telnet session.

(See “Rebooting To Activate Configuration Changes” on page 3-13.)

3. Exit from the terminal program, turn off the terminal, or close the Telnet application program.

Main Menu Features

```
----- CONSOLE - MANAGER MODE -----  
Main Menu  
  
1. Status and Counters ...  
2. Switch Configuration...  
3. Console Passwords...  
4. Event Log  
5. Command Line (CLI)  
6. Reboot Switch  
7. Download OS  
8. Run Setup  
0. Logout  
  
Provides the menu to display configuration, status, and counters.  
To select menu item, press item number, or highlight item and press <Enter>.
```

Figure 3-3. The Main Menu View with Manager Privileges

The Main Menu gives you access to these Menu interface features:

- **Status and Counters:** Provides access to display screens showing switch information, port status and counters, and port and VLAN address tables. (Refer to Appendix B, “Monitoring and Analyzing Switch Operation”.)
- **Switch Configuration:** Provides access to configuration screens for displaying and changing the current configuration settings. (See the Contents listing at the front of this manual.) For a listing of features and parameters configurable through the menu interface, see the “Menu Features List” on page 3-14. For an index of the features covered in the software manuals for your switch, refer to the “Software Feature Index” on page -xxvi.
- **Console Passwords:** Provides access to the screen used to set or change Manager-level and Operator-level passwords, and to delete Manager and Operator password protection. (Refer to the chapter on configuring usernames and passwords in the *Access Security Guide* for your switch.)
- **Event Log:** Enables you to read progress and error messages that are useful for checking and troubleshooting switch operation. (See “Using the Event Log for Troubleshooting Switch Problems” on page C-27.)

- **Command Line (CLI):** Selects the Command Line Interface at the same level (Manager or Operator) that you are accessing in the Menu interface. (Refer to Chapter 4, “Using the Command Line Interface (CLI)”.)
- **Reboot Switch:** Performs a “warm” reboot of the switch, which clears most temporary error conditions, resets the network activity counters to zero, and resets the system up-time to zero. A reboot is required to activate a change in the VLAN Support parameter. (See “Rebooting from the Menu Interface” on page 6-11.)
- **Download OS:** Enables you to download a new switch software version to the switch. (See Appendix A, “File Transfers”.)
- **Run Setup:** Displays the Switch Setup screen for quickly configuring basic switch parameters such as IP addressing, default gateway, logon default interface, and others. (Refer to the *Installation and Getting Started Guide* for your switch.)
- **Logout:** Closes the Menu interface and console session, and disconnects Telnet access to the switch. (See “How to End a Menu Session and Exit from the Console” on page 3-5.)

Screen Structure and Navigation

Menu interface screens include these three elements:

- Parameter fields and/or read-only information such as statistics
- Navigation and configuration actions, such as Save, Edit, and Cancel
- Help line to describe navigation options, individual parameters, and read-only data

For example, in the following System Information screen:

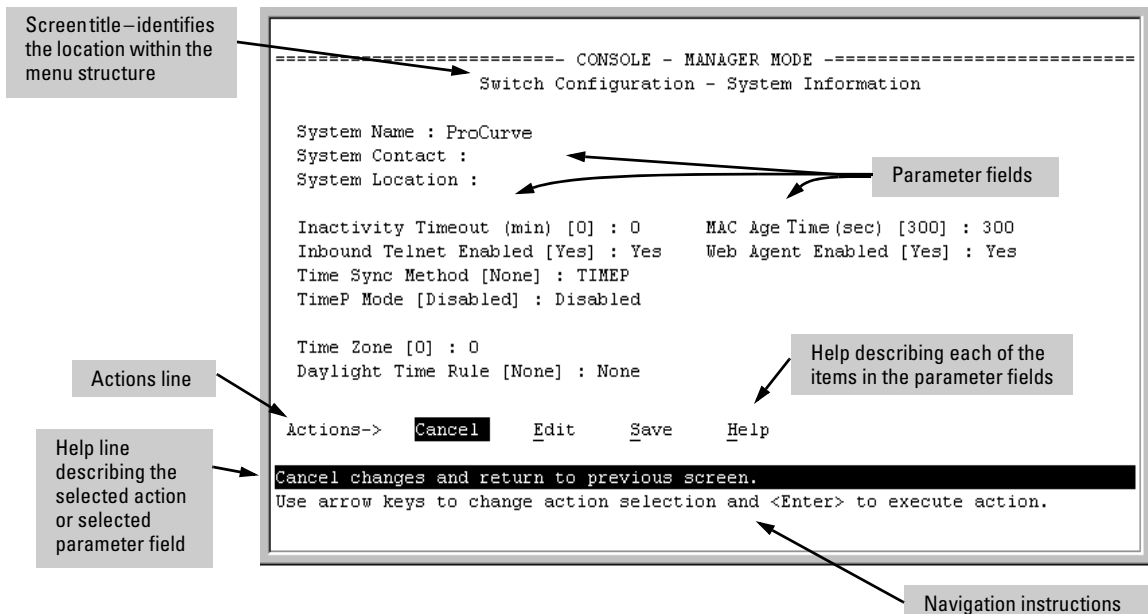


Figure 3-4. Elements of the Screen Structure

“Forms” Design. The configuration screens, in particular, operate similarly to a number of PC applications that use forms for data entry. When you first enter these screens, you see the current configuration for the item you have selected. To change the configuration, the basic operation is to:

1. Press **[E]** to select the **Edit** action.
2. Navigate through the screen making all the necessary configuration changes. (See Table 3-1 on page 3-10.)
3. Press **[Enter]** to return to the **Actions** line. From there you can save the configuration changes or cancel the changes. Cancel returns the configuration to the values you saw when you first entered the screen.

Table 3-1. How To Navigate in the Menu Interface

Task:	Actions:
Execute an action from the “Actions →” list at the bottom of the screen:	<p>Use either of the following methods:</p> <ul style="list-style-type: none"> • Use the arrow keys (←, or →) to highlight the action you want to execute, then press [Enter]. • Press the key corresponding to the capital letter in the action name. For example, in a configuration menu, press [E] to select Edit and begin editing parameter values.
Reconfigure (edit) a parameter setting or a field:	<ol style="list-style-type: none"> 1. Select a configuration item, such as System Name. (See figure 3-4.) 2. Press [E] (for Edit on the Actions line). 3. Use [Tab] or the arrow keys (←, →, ↑, or ↓) to highlight the item or field. 4. Do one of the following: <ul style="list-style-type: none"> – If the parameter has preconfigured values, either use the Space bar to select a new option or type the first part of your selection and the rest of the selection appears automatically. (The help line instructs you to “Select” a value.) – If there are no preconfigured values, type in a value (the Help line instructs you to “Enter” a value). 5. If you want to change another parameter value, return to step 3. 6. If you are finished editing parameters in the displayed screen, press [Enter] to return to the Actions line and do one of the following: <ul style="list-style-type: none"> – To save and activate configuration changes, press [S] (for the Save action). This saves the changes in the startup configuration and also implements the change in the currently running configuration. (See Chapter 6, “Switch Memory and Configuration”.) – To exit from the screen without saving any changes that you have made (or if you have not made changes), press [C] (for the Cancel action). <p>Note: In the menu interface, executing Save activates most parameter changes and saves them in the startup configuration (or flash) memory, and it is therefore not necessary to reboot the switch after making these changes. But if an asterisk appears next to any menu item you reconfigure, the switch will not activate or save the change for that item until you reboot the switch. In this case, rebooting should be done after you have made all desired changes and then returned to the Main Menu.</p> 7. When you finish editing parameters, return to the Main Menu. 8. If necessary, reboot the switch by highlighting Reboot Switch in the Main Menu and pressing [Enter]. (See the Note, above.)
Exit from a read-only screen.	Press [B] (for the B ack action).

To get Help on individual parameter descriptions. In most screens there is a **Help** option in the **Actions** line. Whenever any of the items in the **Actions** line is highlighted, press **[H]**, and a separate help screen is displayed. For example:

----- CONSOLE - MANAGER MODE -----
Switch Configuration - System Information

System Name : ProCurve
System Contact :
System Location :

Inactivity Timeout (min) [0] : 0 MAC Age Time (sec) [300] : 300
Inbound Telnet Enabled [Yes] : Yes Web Agent Enabled [Yes] : Yes
Time Sync Method [None] : TIMEP
TimeP Mode [Disabled] : Disabled

Time Zone [0] : 0
Daylight Time Rule [None] : None

Actions-> **Cancel** Edit Save Help

Cancel changes and return to previous screen.
Use arrow keys to change action selection and <Enter> to execute action.

Annotations:

- Highlight on any item in the Actions line indicates that the Actions line is active.
- The Help line provides a brief descriptor of the highlighted Action item or parameter.
- Pressing **[H]** or highlighting **Help** and pressing **[Enter]** displays Help for the parameters listed in the upper part of the screen

Figure 3-5. Example Showing How To Display Help

To get Help on the actions or data fields in each screen: Use the arrow keys (←, →, ↑, or ↓) to select an action or data field. The help line under the **Actions** items describes the currently selected action or data field.

For guidance on how to navigate in a screen: See the instructions provided at the bottom of the screen, or refer to “Screen Structure and Navigation” on page 3-9.)

Rebooting the Switch

Rebooting the switch from the menu interface

- Terminates all current sessions and performs a reset of the operating system
- Activates any menu interface configuration changes that require a reboot
- Resets statistical counters to zero

(Note that statistical counters can be reset to zero without rebooting the switch.)

To Reboot the switch, use the **Reboot Switch** option in the Main Menu. (Note that **Reboot Switch** is not available if you log on in Operator mode; that is, if you enter an Operator password instead of a manager password at the password prompt.)




Figure 3-6. The Reboot Switch Option in the Main Menu

Rebooting To Activate Configuration Changes. Configuration changes for most parameters in the menu interface become effective as soon as you save them. However, you must reboot the switch in order to implement a change in the **Maximum VLANs to support** parameter. (To access this parameter, go to the Main Menu and select:

2. Switch Configuration

8. VLAN Menu

1. VLAN Support.

If you make configuration changes in the menu interface that require a reboot, the switch displays an asterisk (*) next to the menu item in which the change has been made. For example, if you change and save the value for the **Maximum VLANs to support** parameter, an asterisk appears next to the **VLAN Support** entry in the VLAN Menu screen, and also next to the **Switch Configuration** entry in the Main Menu.

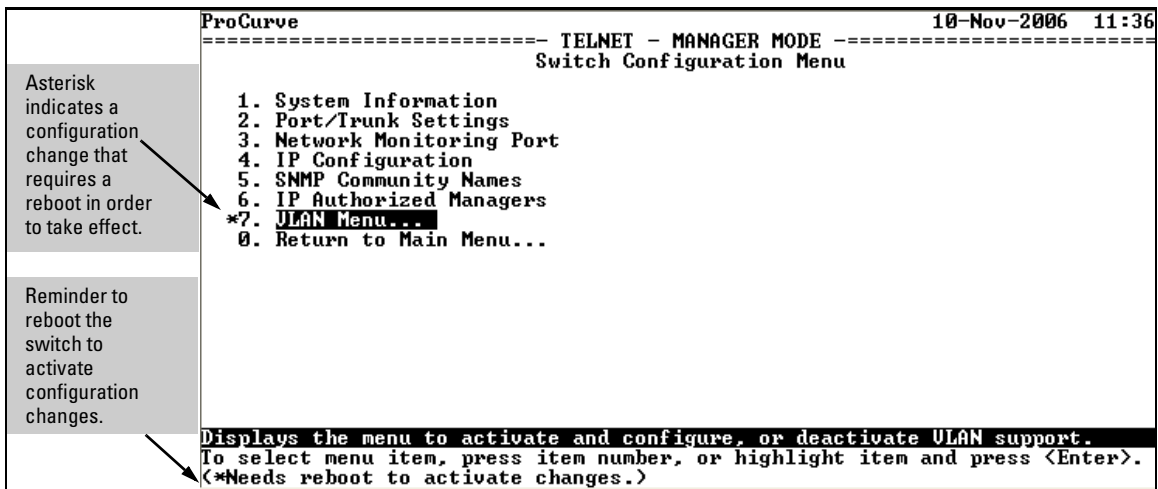


Figure 3-7. Indication of a Configuration Change Requiring a Reboot

To activate changes indicated by the asterisk, go to the Main Menu and select the **Reboot Switch** option.

Note

Executing the **write memory** command in the CLI does not affect pending configuration changes indicated by an asterisk in the menu interface. That is, only a reboot from the menu interface or a **boot** or **reload** command from the CLI will activate a pending configuration change indicated by an asterisk.

Menu Features List

Status and Counters

- General System Information
- Switch Management Address Information
- Port Status
- Port Counters
- Address Table
- Port Address Table

Switch Configuration

- System Information
- Port/Trunk Settings
- Network Monitoring Port
- IP Configuration
- SNMP Community Names
- IP authorized Managers
- VLAN Menu

Console Passwords

Event Log

Command Line (CLI)

Reboot Switch

Download OS (Download Switch Software)

Run Setup

Logout

Where To Go From Here

This chapter provides an overview of the menu interface and how to use it. The following table indicates where to turn for detailed information on how to use the individual features available through the menu interface.

Option:	Turn to:
To use the Run Setup option	Refer to the <i>Installation and Getting Started Guide</i> shipped with the switch.
To view and monitor switch status and counters	Appendix B, "Monitoring and Analyzing Switch Operation"
To learn how to configure and use passwords and other security features	Refer to the <i>Access Security Guide</i> for your switch.
To learn how to use the Event Log	"Using the Event Log for Troubleshooting Switch Problems" on page C-27
To learn how the CLI operates	Chapter 4, "Using the Command Line Interface (CLI)"
To download switch software	Appendix A, "File Transfers"
For a description of how switch memory handles configuration changes	Chapter 6, "Switch Memory and Configuration"
For information on other switch features and how to configure them	Refer to the Feature Index on (page xxvi) at the front of this guide, and to "Sources for More Information" on page 1-4.

Using the Menu Interface
Where To Go From Here

Using the Command Line Interface (CLI)

Contents

Overview	4-2
Accessing the CLI	4-2
Using the CLI	4-2
Privilege Levels at Logon	4-3
Privilege Level Operation	4-4
Operator Privileges	4-4
Manager Privileges	4-5
How To Move Between Levels	4-7
Listing Commands and Command Options	4-8
Listing Commands Available at Any Privilege Level	4-8
Listing Command Options	4-10
Displaying CLI “Help”	4-11
Configuration Commands and the Context Configuration Modes ..	4-13
CLI Control and Editing	4-16
Executing a Prior Command—Redo	4-16
Repeating Execution of a Command	4-16
Using a Command Alias	4-18
CLI Shortcut Keystrokes	4-20

Overview

The CLI is a text-based command interface for configuring and monitoring the switch. The CLI gives you access to the switch's full set of commands while providing the same password protection that is used in the web browser interface and the menu interface.

Accessing the CLI

Like the menu interface, the CLI is accessed through the switch console, and in the switch's factory default state, is the default interface when you start a console session. You can access the console out-of-band by directly connecting a terminal device to the switch, or in-band by using Telnet either from a terminal device or through the web browser interface.

Also, if you are using the menu interface, you can access the CLI by selecting the **Command Line (CLI)** option in the Main Menu.

Using the CLI

The CLI offers these privilege levels to help protect the switch from unauthorized access:

1. Operator
2. Manager
3. Global Configuration
4. Context Configuration

Note

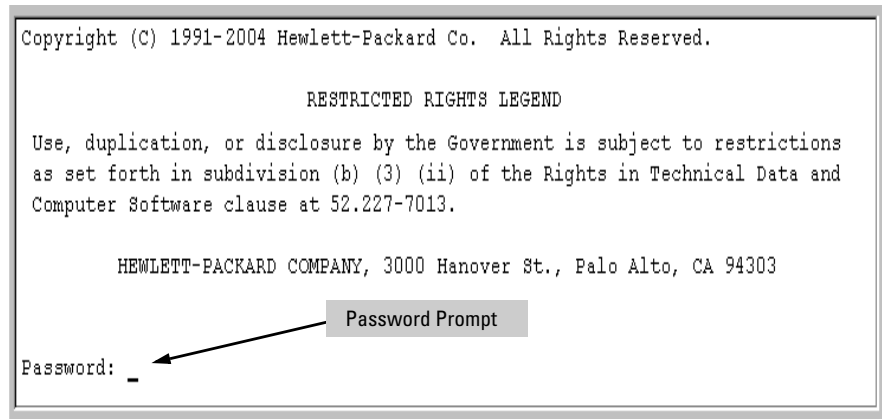
CLI commands are not case-sensitive.

When you use the CLI to make a configuration change, the switch writes the change to the Running-Config file in volatile memory. This allows you to test your configuration changes before making them permanent. To make changes permanent, you must use the **write memory** command to save them to the Startup-Config file in non-volatile memory. If you reboot the switch without first using **write memory**, all changes made since the last reboot or **write memory** (whichever is later) will be lost. For more on switch memory and saving configuration changes, see Chapter 6, “Switch Memory and Configuration”.

Privilege Levels at Logon

Privilege levels control the type of access to the CLI. To implement this control, you must set at least a Manager password. *Without a Manager password configured, anyone having serial port, Telnet, or web browser access to the switch can reach all CLI levels.* (For more on setting passwords, refer to the chapter on usernames and passwords in the *Access Security Guide* for your switch.)

When you use the CLI to log on to the switch, and passwords are set, you will be prompted to enter a password. For example:



```
Copyright (C) 1991-2004 Hewlett-Packard Co. All Rights Reserved.  
  
RESTRICTED RIGHTS LEGEND  
  
Use, duplication, or disclosure by the Government is subject to restrictions  
as set forth in subdivision (b) (3) (ii) of the Rights in Technical Data and  
Computer Software clause at 52.227-7013.  
  
HEWLETT-PACKARD COMPANY, 3000 Hanover St., Palo Alto, CA 94303  
  
Password: _
```

The screenshot shows a terminal window with a grey border. At the top, it displays copyright information for Hewlett-Packard Co. from 1991-2004. Below that is a "RESTRICTED RIGHTS LEGEND" section explaining government use restrictions. The company name and address are listed next. At the bottom, a "Password:" prompt is shown with a single underscore character. A grey callout box labeled "Password Prompt" has an arrow pointing to the underscore.

Figure 4-1. Example of CLI Log-On Screen with Password(s) Set

In the above case, you will enter the CLI at the level corresponding to the password you provide (operator or manager).

If no passwords are set when you log onto the CLI, you will enter at the Manager level. For example:

```
ProCurve# _
```

Caution

ProCurve strongly recommends that you configure a Manager password. If a Manager password is not configured, then the Manager level is not password-protected, and anyone having in-band or out-of-band access to the switch may be able to reach the Manager level and compromise switch and network security. Note that configuring only an Operator password *does not* prevent access to the Manager level by intruders who have the Operator password.

Pressing the Clear button on the front of the switch removes password protection. *For this reason, it is recommended that you protect the switch from physical access by unauthorized persons.* If you are concerned about switch security and operation, you should install the switch in a secure location, such as a locked wiring closet.

Privilege Level Operation




Figure 4-2. Access Sequence for Privilege Levels

Operator Privileges

At the Operator level you can examine the current configuration and move between interfaces without being able to change the configuration. A “>” character delimits the Operator-level prompt. For example:

```
ProCurve> _ (Example of the Operator prompt.)
```

When using **enable** to move to the Manager level, the switch prompts you for the Manager password if one has already been configured.

Manager Privileges

Manager privileges give you three additional levels of access: Manager, Global Configuration, and Context Configuration. A “#” character delimits any Manager prompt. For example:

ProCurve#_ *Example of the Manager prompt.*

- **Manager level:** Provides all Operator level privileges plus the ability to perform system-level actions that do not require saving changes to the system configuration file. The prompt for the Manager level contains only the system name and the “#” delimiter, as shown above. To select this level, enter the **enable** command at the Operator prompt and enter the Manager password, when prompted. For example:

```
ProCurve> enable      Enter enable at the Operator prompt.  
Password:              CLI prompt for the Manager password.  
ProCurve# _            The Manager prompt appears after the  
                         correct Manager password is entered.
```

- **Global Configuration level:** Provides all Operator and Manager level privileges, and enables you to make configuration changes to any of the switch’s software features. The prompt for the Global Configuration level includes the system name and “(config)”. To select this level, enter the **config** command at the Manager prompt. For example:

```
ProCurve# config      Enter config at the Manager prompt.  
ProCurve(config)#_    The Global Config prompt.
```

- **Context Configuration level:** Provides all Operator and Manager privileges, and enables you to make configuration changes in a specific context, such as one or more ports or a VLAN. The prompt for the Context Configuration level includes the system name and the selected context. For example:

```
ProCurve(eth-1)#  
  
ProCurve(vlan-10)#
```

The Context level is useful, for example, for executing several commands directed at the same port or VLAN, or if you want to shorten the command strings for a specific context area. To select this level, enter the specific context at the Global Configuration level prompt. For example, to select the context level for an existing VLAN with the VLAN ID of 10, you would enter the following command and see the indicated result:

```
ProCurve(config)# vlan 10  
  
ProCurve(vlan-10)#
```

Table 4-1. Privilege Level Hierarchy

Privilege Level	Example of Prompt and Permitted Operations		
Operator Privilege			
Operator Level	ProCurve>	show < command > setup	View status and configuration information.
		ping < argument > link-test < argument >	Perform connectivity tests.
		enable	Move from the Operator level to the Manager level.
		menu	Move from the CLI interface to the menu interface.
		logout	Exit from the CLI interface and terminate the console session.
		exit	Terminate the current session (same as logout).
Manager Privilege			
Manager Level	ProCurve#		Perform system-level actions such as system control, monitoring, and diagnostic commands, plus any of the Operator-level commands. For a list of available commands, enter ? at the prompt.
Global Configuration Level	ProCurve(config)#		Execute configuration commands, plus all Operator and Manager commands. For a list of available commands, enter ? at the prompt.
Context Configuration Level	ProCurve(eth-5)# ProCurve(vlan-100)#		Execute context-specific configuration commands, such as a particular VLAN or switch port. This is useful for shortening the command strings you type, and for entering a series of commands for the same context. For a list of available commands, enter ? at the prompt.

How To Move Between Levels

Change in Levels	Example of Prompt, Command, and Result
Operator level <i>to</i> Manager level	ProCurve> enable Password: _ ProCurve# _ After you enter enable , the Password prompt appears. After you enter the Manager password, the system prompt appears with the # symbol:
Manager level <i>to</i> Global configuration level	ProCurve# config ProCurve(config)#
Global configuration level <i>to a</i> Context configuration level	ProCurve(config)# vlan 10 ProCurve(vlan-10)#
Context configuration level <i>to another</i> Context configuration level	ProCurve(vlan-10)# interface e 3 ProCurve(int-3)# The CLI accepts "e" as the abbreviated form of "ethernet".
Move from any level to the preceding level	ProCurve(int-3)# exit ProCurve(config)# exit ProCurve# exit ProCurve>
Move from any level to the Manager level	ProCurve(int-3)# end ProCurve# -or- ProCurve(config)# end ProCurve#

Moving Between the CLI and the Menu Interface. When moving between interfaces, the switch retains the current privilege level (Manager or Operator). That is, if you are at the Operator level in the menu and select the **Command Line Interface (CLI)** option from the Main Menu, the CLI prompt appears at the Operator level.

Changing Parameter Settings. Regardless of which interface is used (CLI, menu interface, or web browser interface), the most recently configured version of a parameter setting overrides any earlier settings for that parameter.

For example, if you use the menu interface to configure an IP address of “X” for VLAN 1 and later use the CLI to configure a different IP address of “Y” for VLAN 1, then “Y” replaces “X” as the IP address for VLAN 1 in the running-config file. If you subsequently execute **write memory** in the CLI, then the switch also stores “Y” as the IP address for VLAN 1 in the startup-config file. (For more on the startup-config and running config files, see Chapter 6, “Switch Memory and Configuration”).

Listing Commands and Command Options

At any privilege level you can:

- List all of the commands available at that level
- List the options for a specific command

Listing Commands Available at Any Privilege Level

At a given privilege level you can list and execute the commands that level offers, plus all of the commands available at preceding levels. For example, at the Operator level, you can list and execute only the Operator level commands. However, at the Manager level, you can list and execute the commands available at both the Operator and Manager levels.

Type “?” To List Available Commands. 1. Typing the ? symbol lists the commands you can execute at the current privilege level. For example, typing ? at the Operator level produces this listing:

```
ProCurve> ?  
  
  enable  
  exit  
  link-test  
  logout  
  menu  
  ping  
  show  
  traceroute  
HPswitch>
```

Figure 4-3. Example of the Operator Level Command Listing

Typing ? at the Manager level produces this listing:

```

ProCurve# ?
boot                Reboot the device.
clear               Clear table/statistics or authorized client public
                   keys.
configure           Enter the Configuration context.
copy                Copy datafiles to/from the switch.
debug              Enable/disable debug logging.
display            Display the running/saved configuration.
end                Return to the Manager Exec context.
erase              Erase the configuration file stored in flash or.
getMIB             Retrieve and display the value of the MIB objects
                   specified.
kill               Kill other active console, telnet, or ssh sessions.
log                Display log events.
page               Toggle paging mode.
print              Execute a command and redirect its output to the device
                   channel for current session.
redo               Re-execute a command from history.
reload             Warm reboot of the switch.
repeat             Repeat execution of a previous command.
setMIB             Set the value of a MIB object.
setup              Enter the 'Switch Setup' screen for basic switch
                   configuration.
-- MORE --, next page: Space, next line: Enter, quit: Control-C

```

When -- MORE -- appears, use the Space bar or [Return] to list additional commands.

Figure 4-4. Example of the Manager-Level Command Listing

When -- **MORE** -- appears, there are more commands in the listing. To list the next screenfull of commands, press the Space bar. To list the remaining commands one-by-one, repeatedly press [Enter].

Typing ? at the Global Configuration level or the Context Configuration level produces similar results.

Use [Tab] To Search for or Complete a Command Word. You can use [Tab] to help you find CLI commands or to quickly complete the current word in a command. To do so, type one or more consecutive characters in a command and then press [Tab] (with no spaces allowed). For example, at the Global Configuration level, if you press [Tab] immediately after typing “t”, the CLI displays the available command options that begin with “t”. For example:

```

ProCurve(config)# t [Tab]
tacacs-server
telnet-server
time
timesync
trunk
telnet
terminal
traceroute
ProCurve(config)# t

```

As mentioned above, if you type part of a command word and press **[Tab]**, the CLI completes the current word (if you have typed enough of the word for the CLI to distinguish it from other possibilities), including hyphenated extensions. For example:

```
ProCurve(config)# port-[Tab]
ProCurve(config)# port-security _
```

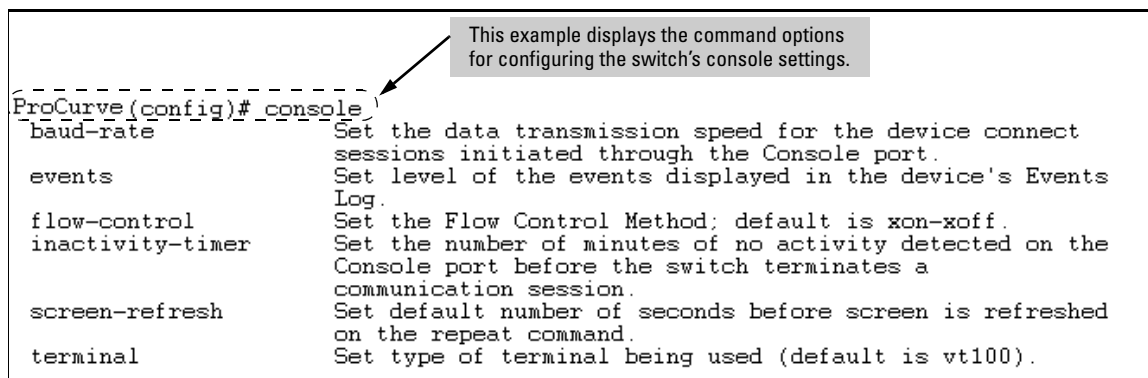
Pressing **[Tab]** after a completed command word lists the further options for that command.

```
ProCurve(config)# qos [Tab]

udp-portSet UDP port based priority.
tcp-portSet TCP port based priority.
device-priorityConfigure device-based priority.
dscp-mapDefine mapping between a DSCP
(Differentiated-Services Codepoint)
value and 802.1p priority.
type-of-serviceConfigure the Type-of-Service
method the device uses to
prioritize IP traffic.
```

Listing Command Options

You can use the CLI to remind you of the options available for a command by entering command keywords followed by **?**. For example, suppose you want to see the command options for configuring the console settings:



The screenshot shows a terminal window with the command `ProCurve(config)# console?` entered. The output lists several options: `baud-rate`, `events`, `flow-control`, `inactivity-timer`, `screen-refresh`, and `terminal`. Each option is followed by a brief description of its function. A grey callout box with an arrow pointing to the command line contains the text: "This example displays the command options for configuring the switch's console settings."

```
ProCurve(config)# console?
  baud-rate      Set the data transmission speed for the device connect
                 sessions initiated through the Console port.
  events         Set level of the events displayed in the device's Events
                 Log.
  flow-control   Set the Flow Control Method: default is xon-xoff.
  inactivity-timer Set the number of minutes of no activity detected on the
                 Console port before the switch terminates a
                 communication session.
  screen-refresh Set default number of seconds before screen is refreshed
                 on the repeat command.
  terminal       Set type of terminal being used (default is vt100).
```

Figure 4-5. Example of How To List the Options for a Specific Command

Displaying CLI “Help”

CLI Help provides two types of context-sensitive information:

- Command list with a brief summary of each command’s purpose
- Detailed information on how to use individual commands

Displaying Command-List Help.

Syntax: help

*Displays a listing of command Help summaries for all commands available at the current privilege level. That is, at the Operator level, executing **help** displays the Help summaries only for Operator-Level commands. At the Manager level, executing **help** displays the Help summaries for both the Operator and Manager levels, and so on.*

For example, to list the Operator-Level commands with their purposes:

```
ProCurve> help
enable          Enter the Manager Exec context.
exit            Return to the previous context or terminate current
               console/telnet session if you are in the Operator
               context level.
link-test       Test the connection to a MAC address on the LAN.
logout          Terminate this console/telnet session.
menu            Change console user interface to menu system.
ping            Send IP Ping requests to a device on the network.
show            Display switch operation information.
traceroute      Send traceroute to a device on the network.
```

Figure 4-6. Example of Context-Sensitive Command-List Help

Displaying Help for an Individual Command.

Syntax: < command-string > help

This option displays Help for any command available at the current context level.

For example, to list the Help for the **interface** command in the Global Configuration privilege level:

```
ProCurve(config)# interface help
Usage: [no] interface [ethernet] PORT-LIST [...]

Description: Enter the Interface Configuration Level, or execute one
             command for that level. Without optional parameters
             specified, the 'interface' command changes the context to
             the Interface Configuration Context Level for execution of
             configuration changes to the port or ports in the PORT-LIST.
             The 'interface [ethernet] PORT-LIST' can be followed by any
             command from the Interface Configuration Context Level in the
             same command line. In this case the context level is not
             changed, but the command is also executed for the port or ports
             in the PORT-LIST. Use 'interface [ethernet] PORT-LIST ?'
             to get a list of all valid commands.
```

Figure 4-7. Example of How To Display Help for a Specific Command

Note that trying to list the help for an individual command from a privilege level that does not include that command results in an error message. For example, trying to list the help for the **interface** command while at the global configuration level produces this result:

```
ProCurve# speed-duplex help
Invalid input: speed-duplex
```

Configuration Commands and the Context Configuration Modes

You can execute any configuration command in the global configuration mode or in selected context modes. However, using a context mode enables you to execute context-specific commands faster, with shorter command strings.

The switch offers interface (port or trunk group) and VLAN context configuration modes:

Port or Trunk-Group Context. Includes port-or-trunk-specific commands that apply only to the selected port(s) or trunk group, plus the global configuration, Manager, and Operator commands. The prompt for this mode includes the identity of the selected port(s):

```
ProCurve(config)# interface c3-c6  
ProCurve(eth-C5-C8)#
```

```
ProCurve(config)# interface trk1  
ProCurve(eth-Trk1)#
```

*Commands executed at configuration level for entering port and **trk1** static trunk-group contexts, and resulting prompts showing port or static trunk contexts..*

```
ProCurve(eth-C5-C8)#  
ProCurve(eth-Trk1)#  
  
ProCurve(eth-C5-C8)# ?  
ProCurve(eth-C5-C8)# ?
```

Lists the commands you can use in the port or static trunk context, plus the Manager, Operator, and context commands you can execute at this level.

Using the Command Line Interface (CLI)

Using the CLI

```
ProCurve(eth-3-6)# ?
| broadcast-limit  Set a broadcast traffic percentage limit.
| disable         Disable port(s).
| enable         Enable port(s).
| flow-control    Enable/disable flow control on the port(s).
| gvrp           Set the GVRP timers on the port (hundreths of a
|               second).
| lacp           Define whether LACP is enabled on the port, and whether
|               it is in active or passive mode when enabled.
| mdix-mode       Set port MDI/MDIX mode (default: auto).
| monitor        Define either the port is to be monitored or not.
| name           Set/unset a name for the port(s).
| qos            Set port-based priority.
| rate-limit     Enable/disable and configure rate-limiting for incoming
|               traffic on the port(s).
| speed-duplex   Define mode of operation for the port(s).
| unknown-vlans  Configure GVRP on the port(s).
|-----|-----|
interface       Enter the Interface Configuration Level, or execute one
|               command for that level.
vlan            Add, delete, edit VLAN configuration or enter a VLAN
|               context.
-- MORE --, next page: Space, next line: Enter, quit: Control-C
```

In the port context, the first block of commands in the "?" listing show the context-specific commands that will affect only ports C3-C6.

The remaining commands in the listing are Manager, Operator, and context commands.

Figure 4-8. Context-Specific Commands Affecting Port Context

VLAN Context . Includes VLAN-specific commands that apply only to the selected VLAN, plus Manager and Operator commands. The prompt for this mode includes the VLAN ID of the selected VLAN. For example, if you had already configured a VLAN with an ID of 100 in the switch:

```
ProCurve(config)# vlan 100
```

Command executed at configuration level to enter VLAN 100 context.

```
ProCurve(vlan-100)#
```

Resulting prompt showing VLAN 100 context.

```
ProCurve(vlan-100)# ?
```

Lists commands you can use in the VLAN context, plus Manager, Operator, and context commands you can execute at this level.




Figure 4-9. Context-Specific Commands Affecting VLAN Context

CLI Control and Editing

Executing a Prior Command—Redo

The redo command executes a prior command in the history list.

Syntax: redo [number | command-str]

Re-executes a command from history. Executes the last command by default.

*number: The position of the command to execute in the history list. When **number** is specified, the n^{th} command starting from the most recent command in the history is executed.*

*command-str: When **command-str** is specified, the most recent command whose name matches the specified string is executed.*

```
ProCurve(config)# show history
2      show arp
1      show flash

ProCurve(config)# redo 2
```

Executes the **show arp** command again.

```
IP ARP table
```

IP Address	MAC Address	Type	Port
-----	-----	-----	-----
15.255.128.1	00000c-07ac00	dynamic	All

Figure 4-10. Example of the redo Command

Repeating Execution of a Command

The **repeat** command executes a previous command in the history list.

Syntax: repeat [cmdlist] [count] [delay]

Repeats execution of a previous command. Repeats the last command by default until a key is pressed.

cmdlist: If a number or range of numbers is specified, the command repeats the n^{th} most recent commands (where “ n ” is the position in the history list).

count: Repeats the command for the number of times specified.

delay: The command repeats execution after a delay for the number of seconds specified.

For example:

```
ProCurve(config)# repeat 1-4,7-8,10 count 2 delay 3
```

```
ProCurve(config)# show history
3      show ver
2      show ip
1      show arp

ProCurve(config)# repeat 1-2
```

Repeats the **show arp** and **show ip** commands.

```
IP ARP table

  IP Address      MAC Address      Type      Port
  -----
15.255.128.1     000000-000000    dynamic

Internet (IP) Service

  IP Routing : Disabled

  Default Gateway :
  Default TTL    : 64
  Arp Age        : 20
  Domain Suffix  :
  DNS server     :

  VLAN           | IP Config  IP Address      Subnet Mask      Proxy ARP
  -----+-----
DEFAULT_VLAN    | DHCP/Bootp 15.255.131.90 255.255.248.0    No No
```

Figure 4-11. Example of repeat Command Using a Range

Using a Command Alias

You can create a simple command alias to use in place of a command name and its options. Choose an alias name that is *not* an existing CLI command already. Existing CLI commands are searched before looking for an alias command; an alias that is identical to an existing command will not be executed.

The **alias** command is executed from the current configuration context (operator, manager, or global). If the command that is aliased has to be executed in the global configuration context, you must execute the alias for that command in the global configuration context as well. This prevents bypassing the security in place for a particular context.

ProCurve recommends that you configure no more than 128 aliases.

Syntax: [no] alias <name> <command>

*Creates a shortcut alias name to use in place of a commonly used command. The **alias** command is executed from the current config context.*

name: *Specifies the new command name to use to simplify keystrokes and aid memory.*

command: *Specifies an existing command to be aliased. The command must be enclosed in quotes.*

*Use the **no** form of the command to remove the alias.*

For example, if you use the **show interface custom** command to specify the output, you can configure an alias for the command to simplify execution. It is recommended that you use an alias that does not have an existing tab completion in the CLI. For example, using an alias that starts with “show” or “int” would complete to “show” and “interface” respectively when you use the tab completion function.

```
ProCurve(config)# show int custom 1-4 port name:4 type vlan intrusion speed
enabled mdi
```

```
Status and Counters - Custom Port Status
```

Port Name	Type	VLAN	Intrusion Alert	Speed	Enabled	MDI-mode
1	Acco	100/1000T	1	No	1000FDx	Yes Auto
2	Huma	100/1000T	1	No	1000FDx	Yes Auto
3	Deve	100/1000T	1	No	1000FDx	Yes Auto
4	Labl	100/1000T	1	No	1000FDx	Yes Auto

```
ProCurve(config)# alias sic "show int custom 1-4 port name:4 type vlan intrusion
speed enabled mdi"
```

```
ProCurve(config)#
```

```
ProCurve(config)# sic
```

```
Status and Counters - Custom Port Status
```

Port Name	Type	VLAN	Intrusion Alert	Speed	Enabled	MDI-mode
1	Acco	100/1000T	1	No	1000FDx	Yes Auto
2	Huma	100/1000T	1	No	1000FDx	Yes Auto
3	Deve	100/1000T	1	No	1000FDx	Yes Auto
4	Labl	100/1000T	1	No	1000FDx	Yes Auto

Figure 4-12. Example of Using the Alias Command with show int custom

Note

Remember to enclose the command being aliased in quotes.

Command parameters for the aliased command can be added at the end of the alias command string. For example:

```
ProCurve(config)# alias sc "show config"
ProCurve(config)# sc status
```

To change the command that is aliased, re-execute the alias name with new command options. The new options are used when the alias is executed.

To display the alias commands that have been configured, enter the **show alias** command.

```
ProCurve(config)# show alias

      Name                Command
-----
sc                show config
sic               show int custom 1-4 port name:4 type vlan intrusion
                  speed enabled mdi
```

Figure 4-13. Example of Alias Commands and Their Configurations

CLI Shortcut Keystrokes

Keystrokes	Function
[Ctrl] [A]	Jumps to the first character of the command line.
[Ctrl] [B] or ←	Moves the cursor back one character.
[Ctrl] [C]	Terminates a task and displays the command prompt.
[Ctrl] [D]	Deletes the character at the cursor.
[Ctrl] [E]	Jumps to the end of the current command line.
[Ctrl] [F] or →	Moves the cursor forward one character.
[Ctrl] [K]	Deletes from the cursor to the end of the command line.
[Ctrl] [L] or [Ctrl] [R]	Repeats current command line on a new line.
[Ctrl] [N] or ↓	Enters the next command line in the history buffer.
[Ctrl] [P] or ↑	Enters the previous command line in the history buffer.
[Ctrl] [U] or [Ctrl] [X]	Deletes from the cursor to the beginning of the command line.
[Ctrl] [W]	Deletes the last word typed.
[Esc] [B]	Moves the cursor backward one word.
[Esc] [D]	Deletes from the cursor to the end of the word.
[Esc] [F]	Moves the cursor forward one word.
[Backspace]	Deletes the first character to the left of the cursor in the command line.
[Spacebar]	Moves the cursor forward one character.

Using the ProCurve Web Browser Interface

Contents

Overview	5-2
General Features	5-3
Starting a Web Browser	
Interface Session with the Switch	5-4
Using a Standalone Web Browser in a PC or UNIX Workstation	5-4
Using ProCurve Manager (PCM) or ProCurve Manager Plus (PCM+)	5-5
Tasks for Your First ProCurve Web Browser Interface Session . .	5-7
Viewing the “First Time Install” Window	5-7
Security: Creating Usernames and Passwords in the Browser Interface	5-8
Entering a User Name and Password	5-10
Using a User Name	5-10
If You Lose the Password	5-10
Online Help for the Web Browser Interface	5-11
Support/Mgmt URLs Feature	5-12
Support URL	5-13
Help and the Management Server URL	5-13
Using the PCM Server for Switch Web Help	5-14
Status Reporting Features	5-16
The Overview Window	5-16
The Port Utilization and Status Displays	5-17
Port Utilization	5-17
Port Status	5-19
The Alert Log	5-20
Sorting the Alert Log Entries	5-20
Alert Types and Detailed Views	5-21

Using the ProCurve Web Browser Interface
Contents

Status Indicators	5-22
Setting Fault Detection Policy	5-23

Overview

The ProCurve web browser interface built into the switch lets you easily access the switch from a browser-based PC on your network. This lets you do the following:

- Optimize your network uptime by using the Alert Log and other diagnostic tools
- Make configuration changes to the switch
- Maintain security by configuring usernames and passwords

This chapter covers the following:

- General features (page 5-4).
- Starting a web browser interface session (page 5-5)
- Tasks for your first web browser interface session (page 5-8):
 - Creating usernames and passwords in the web browser interface (page 5-9)
 - Selecting the fault detection configuration for the Alert Log operation (page 5-24)
 - Getting access to online help for the web browser interface (page 5-12)
- Description of the web browser interface:
 - Overview window and tabs (page 5-17)
 - Port Utilization and Status displays (page 5-18)
 - Alert Log and Alert types (page 5-21)
 - Setting the Fault Detection Policy (page 5-24)

Note

You can disable access to the web browser interface by either executing **no web-management** at the Command Prompt or changing the **Web Agent Enabled** parameter setting to **No** (page 7-4).

For information on operating system, browser, and Java versions for the switches covered in this guide, go to the ProCurve Networking web site at www.procurve.com/faqs, select your switch (for example, **ProCurve Switch 8212zl**), and then scroll to **General Product Information**.

General Features

The web browser interface includes these features:

Switch Identity and Status:

- General system data
- Software version
- Redundant Management Module software version
- IP address
- Status Overview
- Port utilization
- Port counters
- Port status
- Redundancy Status
- Alert log

Switch Configuration:

- Device view
- Port configuration
- VLAN configuration
- Fault detection
- Quality of service (QoS)
- Port monitoring (mirroring)
- System information
- IP configuration
- Support and management server URLs
- Device features (Spanning Tree On/Off, VLAN selection, and IGMP)

Switch Security:

- User names and passwords
- Authorized Addresses
- Intrusion Log
- SSL
- RADIUS authentication (Refer to the *Access Security Guide*.)

Switch Diagnostics:

- Ping/Link Test
- Device reset
- Configuration report

Starting a Web Browser Interface Session with the Switch

You can start a web browser session in the following ways:

- Using a standalone web browser on a network connection from a PC or UNIX workstation:
 - Directly connected to your network
 - Connected through remote access to your network
- Using a network management station running ProCurve Manager on your network

Using a Standalone Web Browser in a PC or UNIX Workstation

This procedure assumes that you are using a compatible web browser and that the switch is configured with an IP address accessible from your PC or workstation. (For more on assigning an IP address, refer to “IP Configuration” on page 8-2.)

1. Ensure that the Java™ applets are enabled for your browser. For more information on this topic, refer to your browser’s online Help.
2. Use the web browser to access the switch. If your network includes a Domain Name Server (DNS), your switch’s IP address may have a name associated with it (for example, **switch8212**) that you can type in the **Location or Address** field instead of the IP address. Using DNS names typically improves browser performance. Contact your network administrator to enquire about DNS names associated with your ProCurve switch.

Type the IP address (or DNS name) of the switch in the browser **Location or Address** (URL) field and press **[Enter]**. (It is not necessary to include **http://**.)

switch5308 **[Enter]** (example of a DNS-type name)

10.11.12.195 **[Enter]** (example of an IP address)

Using ProCurve Manager (PCM) or ProCurve Manager Plus (PCM+)

ProCurve Manager and ProCurve Manager Plus are designed for installation on a network management workstation. For this reason, the system requirements are different from the system requirements for accessing the switch's web browser interface from a non-management PC or workstation. For PCM and PCM+ requirements, refer to the information provided with the software.

This procedure assumes that:

- You have installed the recommended web browser on a PC or workstation that serves as your network management station.
- The networked device you want to access has been assigned an IP address and (optionally) a DNS name, and has been discovered by PCM or PCM+. (For more on assigning an IP address, refer to “IP Configuration” on page 8-2.)

To establish a web browser session with PCM or PCM+ running, do the following on the network management station:

1. Make sure the Java™ applets are enabled for your web browser. If they are not, refer to the web browser online Help for specific information on enabling the Java applets.
2. In the **Interconnected Devices** listing under **Network Manager Home** (in the PCM/PCM+ sidebar), right-click on the model number of the device you want to access.
3. The web browser interface automatically starts with the Status Overview window displayed for the selected device, as shown in Figure 5-1.

Note

If the Registration window appears, click on the **Status** tab.

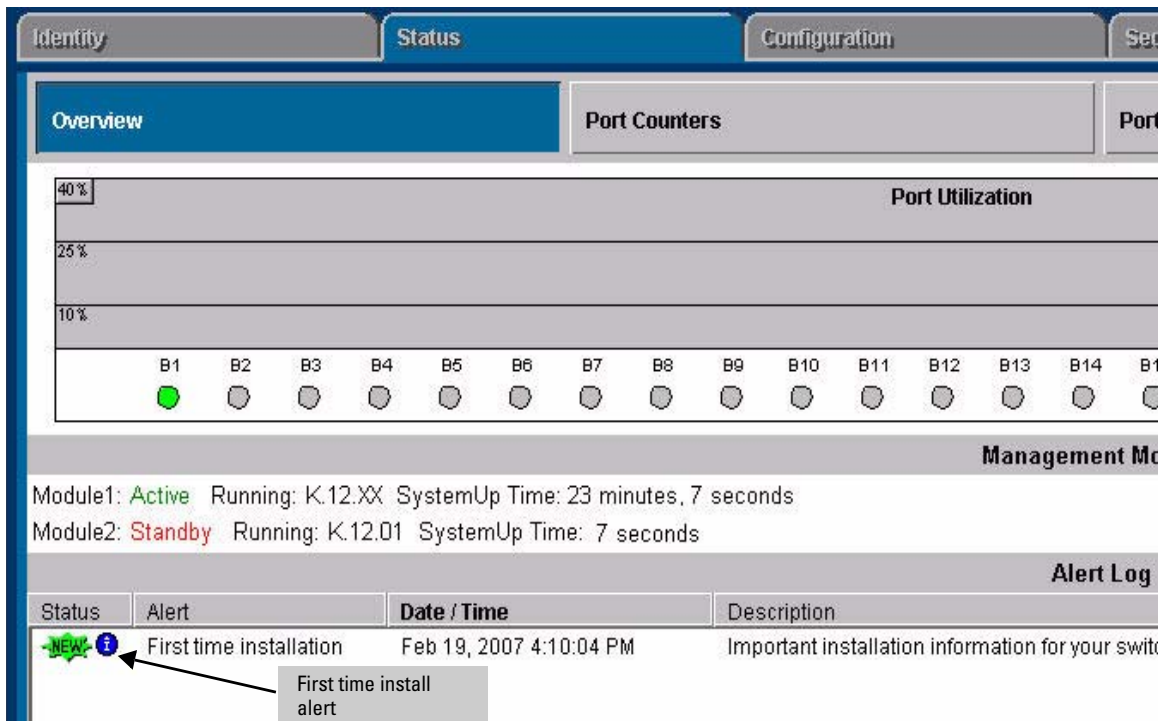


Figure 5-1. Example of Status Overview Screen

Tasks for Your First ProCurve Web Browser Interface Session

The first time you access the web browser interface, there are three tasks you should perform:

- Review the “First Time Install” window
- Set Manager and Operator passwords
- Set access to the web browser interface online help

Viewing the “First Time Install” Window

When you access the switch’s web browser interface for the first time, the Alert log contains a “First Time Install” alert, as shown in figure 5-2. This gives you information about first time installations, and provides an immediate opportunity to set passwords for security and to specify a Fault Detection policy, which determines the types of messages that will be displayed in the Alert Log.

Double click on **First Time Install** in the Alert log (figure 5-1 on page 5-7). The web browser interface then displays the “First Time Install” window, below.

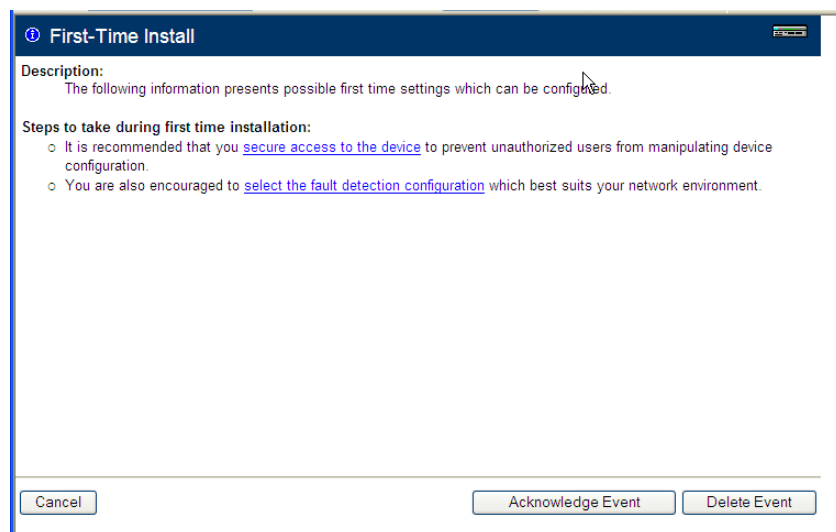


Figure 5-2. First-Time Install Window

This window is the launching point for the basic configuration you need to perform to set web browser interface passwords for maintaining security and a fault detection policy, which determines the types of messages that the Alert Log displays.

To set web browser interface passwords, click on **secure access to the device** to display the Device Passwords screen, and then go to the next page. (You can also access the password screen by clicking on the **Security** tab.)

To set Fault Detection policy, click on **select the fault detection configuration** in the second bullet in the window and go to the section, “Setting Fault Detection Policy” on page 5-24. (You can also access the password screen by clicking on the **Configuration** tab, and then the **[Fault Detection]** key.)

Security: Creating Usernames and Passwords in the Browser Interface

Note

On the switches covered in this guide you can also configure RADIUS authentication for web browser interface access. For more information, refer to the chapter titled “RADIUS Authentication and Accounting” in the *Access Security Guide* for your switch.

You may want to create both a username and a password to create access security for your switch. There are two levels of access to the interface that can be controlled by setting user names and passwords:

- **Operator Setting.** An Operator-level user name and password allows read-only access to most of the web browser interface, but prevents access to the Security window.
- **Manager Setting.** A Manager-level user name and password allows full read/write access to the web browser interface.

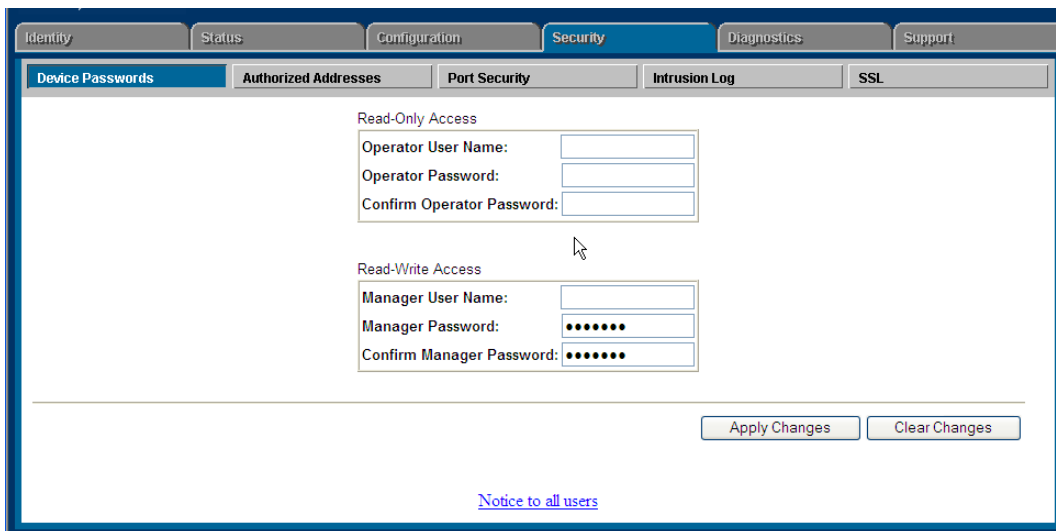


Figure 5-3. The Device Passwords Window

To set the passwords:

1. Access the Device Passwords screen by one of the following methods:
 - If the Alert Log includes a “First Time Install” event entry, double click on this event, then, in the resulting display, click on the **secure access to the device** link.
 - Select the **Security** tab.
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]** to activate the user names and passwords.

Note

Passwords you assign in the web browser interface will overwrite previous passwords assigned in either the web browser interface, the CLI, or the menu interface. That is, the most recently assigned passwords are the switch's passwords, regardless of which interface was used to assign the string.

Entering a User Name and Password




Figure 5-4. Example of the Password Prompt in the Web Browser Interface

The manager and operator passwords are used to control access to all switch interfaces. Once set, you will be prompted to supply the password every time you try to access the switch through any of its interfaces. The password you enter determines the capability you have during that session:

- Entering the manager password gives you full read/write/troubleshooting capabilities
- Entering the operator password gives you read and limited troubleshooting capabilities.

Using a User Name

If you also set user names in the web browser interface screen, you must supply the correct user name for web browser interface access. If a user name has not been set, then leave the User Name field in the password window blank.

Note that the Command Prompt and switch console interfaces use only the password, and do not prompt you for the User Name.

If You Lose the Password

If you lose the passwords, you can clear them by pressing the Clear button on the front of the switch. *This action deletes all password and user name protection from all of the switch's interfaces.*

The Clear button is provided for your convenience, but its presence means that if you are concerned with the security of the switch configuration and operation, you should make sure the switch is installed in a secure location, such as a locked wiring closet. (For more information, refer to “Front Panel Security” in the chapter titled “Configuring Username and Password Security” in the Access Security Guide for your switch.)

Online Help for the Web Browser Interface

Online Help is available for the web browser interface. You can use it by clicking on the “Help” text in the lower right corner of any of the web browser interface screens.




Figure 5-5. The Help Button

Context-sensitive help is provided for the screen you are on.

Note

To access the online Help for the ProCurve web browser interface, you need either ProCurve Manager (version 1.5 or greater) installed on your network or an active connection to the World Wide Web. Otherwise, Online help for the web browser interface will not be available.

For more on Help access and operation, refer to “Help and the Management Server URL” on page 5-14.

Support/Mgmt URLs Feature

The Support/Mgmt URLs window enables you to change the World Wide Web Universal Resource Locator (URL) for two functions:

- **Support URL** – A support information site for your switch
- **Management Server URL** – The web site for web browser online Help

1. Click Here

2. Click Here

ProCurve Networking
HP Innovation

procurve - Status: Information
ProCurve Switch

Identity Status Configuration Security Diagnostics Support

Device View Fault Detection System Info IP Configuration
Port Configuration Quality of Service Monitor Port Device Features
VLAN Configuration Support/Mgmt URL PoE Configuration

Support URL:

Management Server URL:

Apply Changes Clear Changes

[Notice to all users](#)

3. Enter one of the following (or use the default setting):

- The URL for the support information source you want the switch to access when you click on the web browser interface Support tab. The default is the URL for the ProCurve Networking home page.
- The URL of a PCM (ProCurve Network Manager) workstation or other server for the online Help files for this web browser interface. (The default setting accesses the switch's browser-based Help on the ProCurve World Wide Web site.) Note that if you install PCM in your network, the PCM management station acts as the web browser Help server and automatically inserts the necessary URL in this field.)

4. Click on **Apply Changes**

Figure 5-6. The Default Support/Mgmt URLs Window

Support URL

This is the site the switch accesses when you click on the **Support** tab on the web browser interface. The default URL is:

www.procurve.com

which is the World Wide Web site for ProCurve networking products. Click on **technical support** on that page to get support information regarding your switch, including white papers, software updates, and more.

As an alternative, you can replace the ProCurve URL with the URL for a local site used for logging reports on network performance or other support activities.

Help and the Management Server URL

The **Management Server URL** field specifies the URL the switch uses to find online Help for the web browser interface.

- If you install PCM (ProCurve Manager) in your network, the PCM management station acts as the web browser Help server for the switch and automatically inserts the necessary URL in this field. For more on the option, see “Using the PCM Server for Switch Web Help” on page 5-15.)
- In the default configuration (and if PCM is not running on your network) this field is set to the URL for accessing online Help from the ProCurve Networking web site:

www.hp.com/rnd/device_help

Using this option, the Help files are automatically available if your workstation can access the World Wide Web. In this case, if Online Help fails to operate, ensure that the above URL appears in the **Management Server URL** field shown in Figure 5-7:

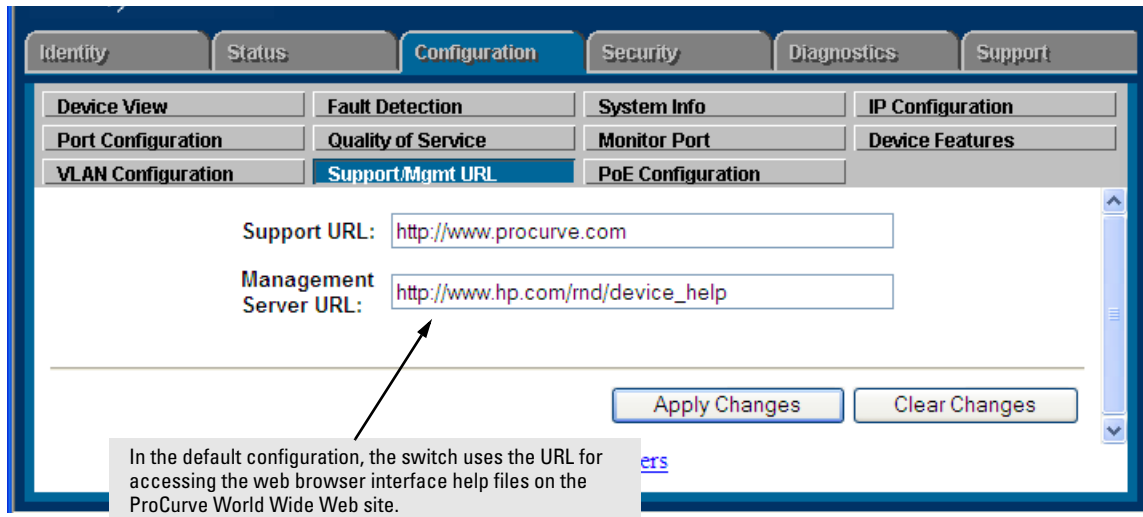


Figure 5-7. How To Access Web Browser Interface Online Help

Using the PCM Server for Switch Web Help

For ProCurve devices that support the “Web Help” feature, you can use the PCM server to host the switch help files for devices that do not have HTTP access to the ProCurve Support Web site.

1. Go to the ProCurve Support web site to get the Device Help files:

www.hp.com/rnd/device_help/

2. Copy the Web help files to the PCM server, under:

```
C:\program files\hewlett-packard\pnm\server\webroot\  
rnd\sevice_help\help\hpwnd\webhelp
```

3. Add an entry, or edit the existing entry in the Discovery portion of the global properties (globalprops.prp) in PCM to redirect the switches to the help files on the PCM server. For example:

```
Global {  
  TempDir=data/temp  
  ...  
  Discovery{  
    ...  
    ...  
    DeviceHelpUrlRedirect=http://15.29.37.12.8040/rnd/device_help  
    ...  
  }  
}
```

You will enter the IP address for your PCM server. 8040 is the standard port number to use.

4. Restart the Discovery process for the change to be applied.

Note

Changing the Discovery's Global properties file will redirect the Device Help URL for all devices.

If you just want to change the Device Help URL for a particular device, then go to the Configuration tab on the Web UI for that device and select the "Support/Mgmt URL" button. Edit the entry in the "Management Server URL" field for the device to point to the PCM server; for example:

http://15.29.37.12.8040/rnd/device_help

Status Reporting Features

Browser elements covered in this section include:

- The Overview window (below)
- Port utilization and status (page 5-18)
- The Alert log (page 5-21)
- The Status bar (page 5-23)

The Overview Window

The Overview Window is the home screen for any entry into the web browser interface. The following figure identifies the various parts of the screen.

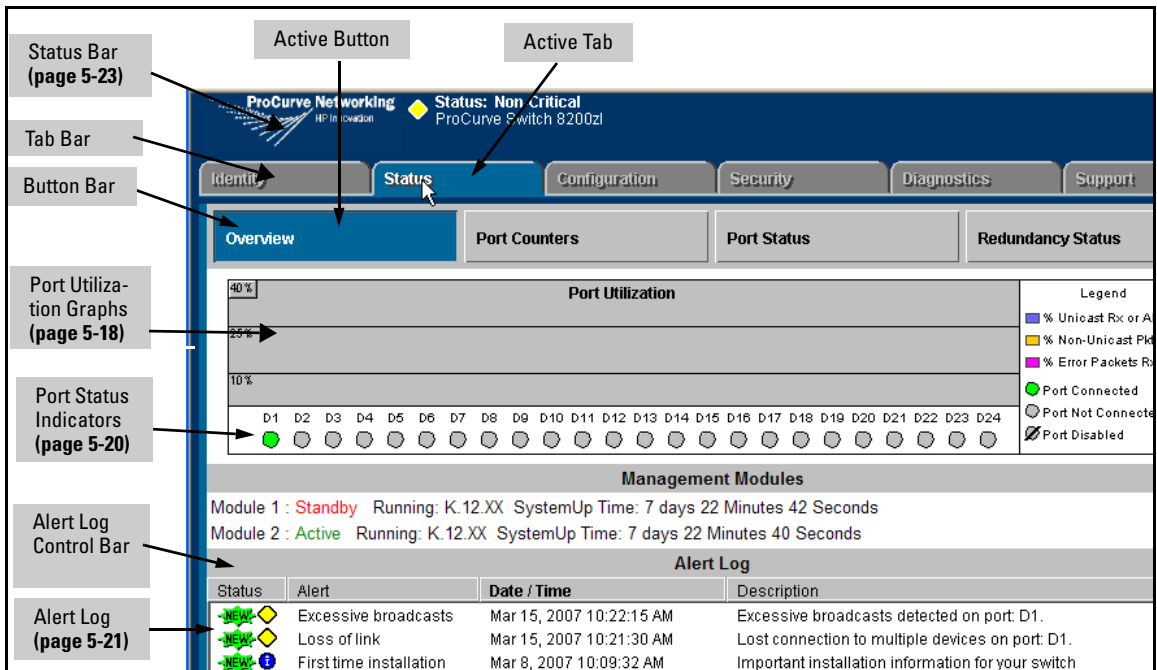


Figure 5-8. The Status Overview Window

Policy Management and Configuration. PCM can perform network-wide policy management and configuration of your switch. The Management Server URL field (page 5-14) shows the URL for the management station performing that function. For more information, refer to the documentation provided with the PCM software.

The Port Utilization and Status Displays

The Port Utilization and Status displays show an overview of the status of the switch and the amount of network activity on each port. The following figure shows a sample reading of the Port Utilization and Port Status.

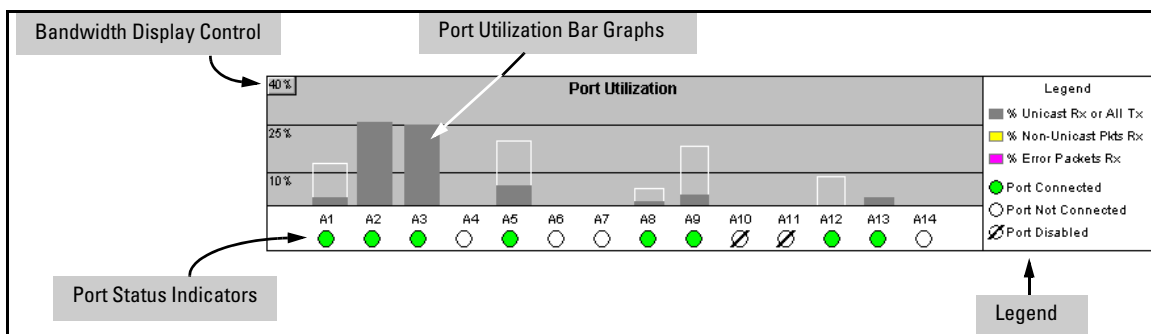


Figure 5-9. The Graphs Area

Port Utilization

The Port Utilization bar graphs show the network traffic on the port with a breakdown of the packet types that have been detected (unicast packets, non-unicast packets, and error packets). The Legend identifies traffic types and their associated colors on the bar graph:

- **% Unicast Rx & All Tx:** This is all unicast traffic received and all transmitted traffic of any type. This indicator (a blue color on many systems) can signify either transmitted or received traffic.
- **% Non-Unicast Pkts Rx:** All multicast and broadcast traffic received by the port. This indicator (a gold color on many systems) enables you to know “at-a-glance” the source of any non-unicast traffic that is causing high utilization of the switch. For example, if one port is receiving heavy broadcast or multicast traffic, all ports will become highly utilized. By color-coding the received broadcast and multicast utilization, the bar graph quickly and easily identifies the offending port. This makes it faster and easier to discover the exact source of the heavy traffic because you don’t have to examine port counter data from several ports.

- **% Error Pkts Rx:** All error packets received by the port. (This indicator is a reddish color on many systems.) Although errors received on a port are not propagated to the rest of the network, a consistently high number of errors on a specific port may indicate a problem on the device or network segment connected to the indicated port.
- **Maximum Activity Indicator:** As the bars in the graph area change height to reflect the level of network activity on the corresponding port, they leave an outline to identify the maximum activity level that has been observed on the port.

Utilization Guideline. A network utilization of 40% is considered the maximum that a typical Ethernet-type network can experience before encountering performance difficulties. If you observe utilization that is consistently higher than 40% on any port, click on the Port Counters button to get a detailed set of counters for the port.

To change the amount of bandwidth the Port Utilization bar graph shows. Click on the bandwidth display control button in the upper left corner of the graph. (The button shows the current scale setting, such as 40%.) In the resulting menu, select the bandwidth scale you want the graph to show (3%, 10%, 25%, 40%, 75%, or 100%), as shown in figure figure 5-10.

Note that when viewing activity on a gigabit port, you may want to select a lower value (such as 3% or 10%). This is because the bandwidth utilization of current network applications on gigabit links is typically minimal, and may not appear on the graph if the scale is set to show high bandwidth utilization.




Figure 5-10. Changing the Graph Area Scale

To display values for each graph bar. Hold the mouse cursor over any of the bars in the graph, and a pop-up display is activated showing the port identification and numerical values for each of the sections of the bar, as shown in figure 5-11 (next).




Figure 5-11. Display of Numerical Values for the Bar

Port Status

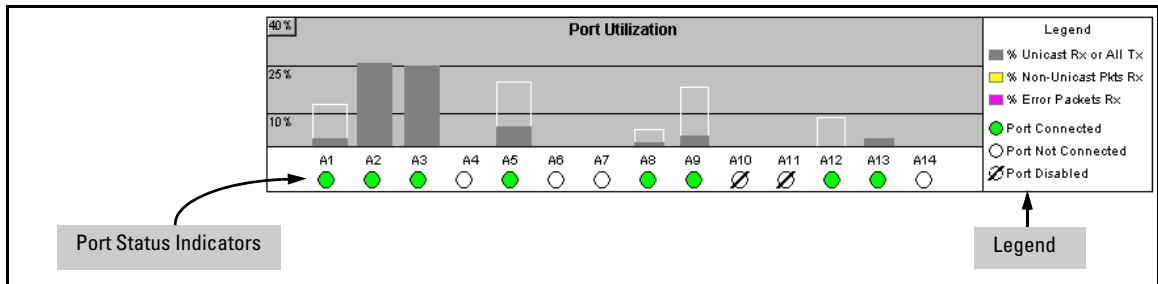


Figure 5-12. The Port Status Indicators and Legend

The Port Status indicators show a symbol for each port that indicates the general status of the port. There are four possible statuses:

- **Port Connected** – the port is enabled and is properly connected to an active network device.
- **Port Not Connected** – the port is enabled but is not connected to an active network device. A cable may not be connected to the port, or the device at the other end may be powered off or inoperable, or the cable or connected device could be faulty.
- **Port Disabled** – the port has been configured as disabled through the web browser interface, the switch console, or SNMP network management.
- **Port Fault-Disabled** – a fault condition has occurred on the port that has caused it to be auto-disabled. Note that the Port Fault-Disabled symbol will be displayed in the legend only if one or more of the ports is in that status. See Appendix B, “Monitoring and Analyzing Switch Operation” for more information.

The Alert Log

The web browser interface Alert Log, shown in the lower half of the screen, shows a list of network occurrences, or *alerts*, that were detected by the switch. Typical alerts are **Broadcast Storm**, indicating an excessive number of broadcasts received on a port, and **Problem Cable**, indicating a faulty cable. A full list of alerts is shown in the table on page 5-22.

Status	Alert	Date / Time	Description
	First time installation	Oct 3, 2005 11:02:47 AM	Important installation information for your switch
	Excessive CRC/ alignment errors	Oct 3, 2005 1:39:03 PM	Excessive CRC/Alignment errors on port: A1.
	Excessive broadcasts	Oct 3, 2005 1:39:03 PM	Excessive broadcasts detected on port: A1.
	Loss of link	Oct 3, 2005 1:38:28 PM	Lost connection to multiple devices on port: A1.

Refresh Open Event Acknowledge Selected Events Delete Selected Events

[Notice to all users](#)

Figure 5-13. Example of the Alert Log

Each alert has the following fields of information:

- **Status** – The level of severity of the event generated. Severity levels can be Information, Normal, Warning, and Critical. If the alert is new (has not yet been acknowledged), the New symbol is also in the Status column.
- **Alert** – The specific event identification.
- **Date/Time** – The date and time the event was received by the web browser interface. This value is shown in the format: **DD-MM-YY HH:MM:SS AM/PM**, for example, **16-Sep-99 7:58:44 AM**.
- **Description** – A short narrative statement that describes the event. For example, **Excessive CRC/Alignment errors on port: 8**.

Sorting the Alert Log Entries

The alerts are sorted, by default, by the Date/Time field with the most recent alert listed at the top of the list. The second most recent alert is displayed below the top alert and so on. If alerts occurred at the same time, the simultaneous alerts are sorted by order in which they appear in the MIB.

Bold characters in a column heading indicate that the alert field alert log entries. You can sort by any of the other columns by clicking on the column heading. The **Alert** and **Description** columns are sorted alphabetically, while the **Status** column is sorted by severity type, with more critical severity indicators appearing above less critical indicators.

Alert Types and Detailed Views

As of June, 2007, the web browser interface generates the following alert types:

- Auto Partition
- Backup Transition
- Excessive broadcasts
- Excessive CRC/alignment errors
- Excessive jabbering
- Excessive late collisions
- First Time Install
- Full-Duplex Mismatch
- Half-Duplex Mismatch
- High collision or drop rate
- Loss of Link
- Mis-Configured SQE
- Network Loop
- Polarity Reversal
- Security Violation
- Stuck 10BaseT Port
- Too many undersized (runt)/giant packets
- Transceiver Hot Swap

Note

When troubleshooting the sources of alerts, it may be helpful to check the switch's Port Status and Port Counter windows, or use the CLI or menu interface to view the switch's Event Log.

When you double click on an Alert Entry, the web browser interface displays a separate window showing information about the event. This view includes a description of the problem and a possible solution. It also provides three management buttons:

- **Acknowledge Event** – removes the New symbol from the log entry
- **Delete Event** – removes the alert from the Alert Log
- **Cancel** – closes the detail view with no change to the status of the alert and returns you to the Overview screen.

For example, figure 5-14 shows a sample detail view describing an Excessive CRC/Alignment Error alert.

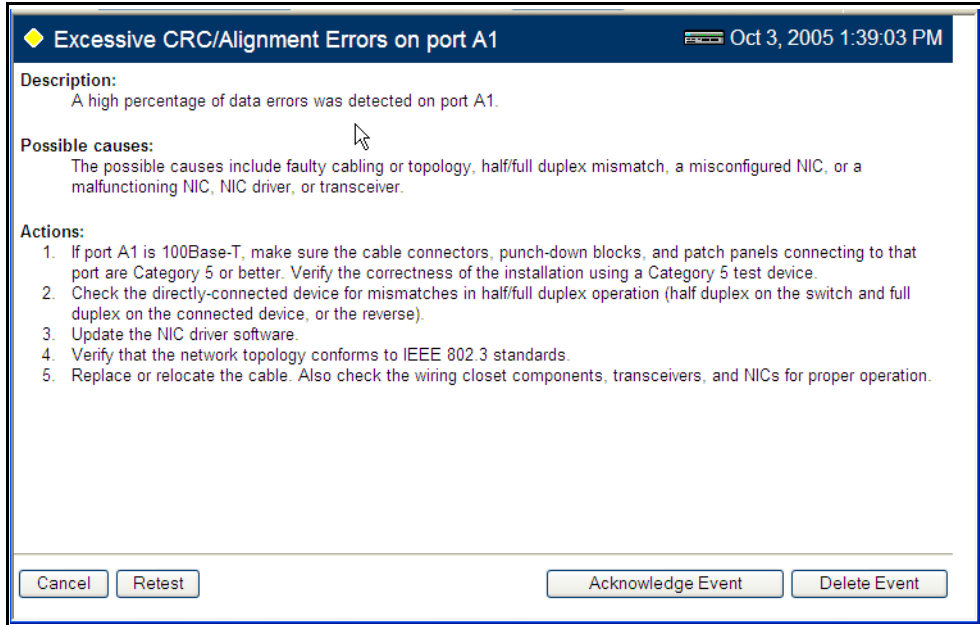


Figure 5-14. Example of Alert Log Detail View

Status Indicators

The status indicators use icons to show the severity of alerts in the current display of the Alert Log. This indicator can be one of four shapes and colors, as shown below.

Table 5-1. Status Indicator Key

Color	Switch Status	Status Indicator Shape
Blue	Normal Activity; "First time installation" information available in the Alert log.	
Green	Normal Activity	
Yellow	Warning	
Red	Critical	

Setting Fault Detection Policy

One of the powerful features in the web browser interface is the Fault Detection facility. For your switch, this feature controls the types of alerts reported to the Alert Log based on their level of severity.

Set this policy in the Fault Detection window (figure 5-15).

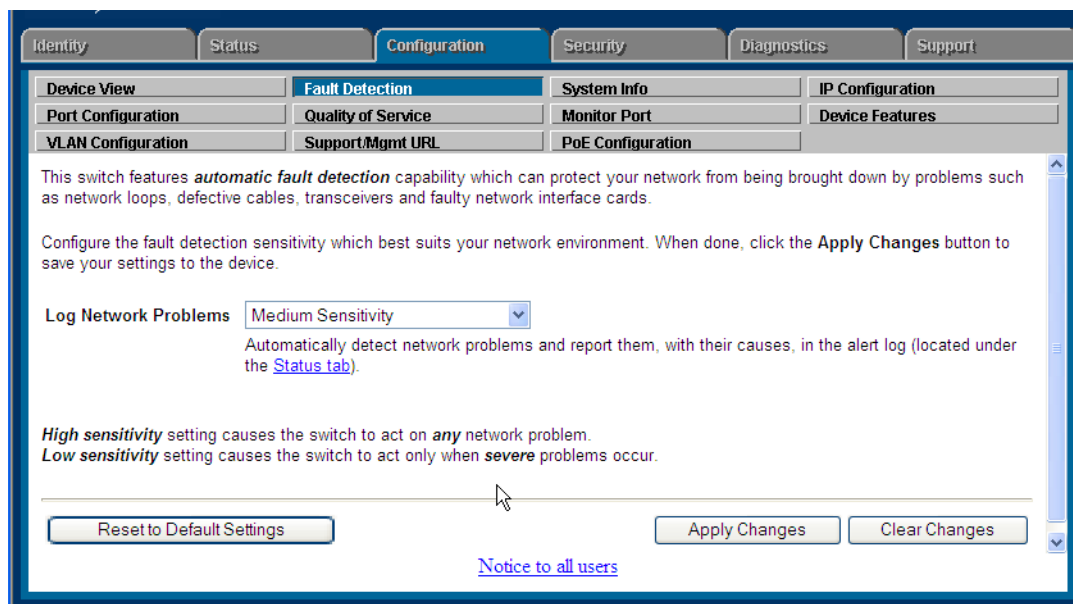


Figure 5-15. The Fault Detection Window

The Fault Detection screen contains a list box for setting fault detection and response policy, and enables you to set the sensitivity level at which a network problem should generate an alert and send it to the Alert Log.

To provide the most information on network problems in the Alert Log, the recommended sensitivity level for **Log Network Problems** is **High Sensitivity**. The Fault Detection settings are:

- **High Sensitivity.** This policy directs the switch to send all alerts to the Alert Log. This setting is most effective on networks that have none or few problems.
- **Medium Sensitivity.** This policy directs the switch to send alerts related to network problems to the Alert Log. If you want to be notified of problems which cause a noticeable slowdown on the network, use this setting.
- **Low Sensitivity.** This policy directs the switch to send only the most severe alerts to the Alert Log. This policy is most effective on a network where there are normally a lot of problems and you want to be informed of only the most severe ones.
- **Never.** Disables the Alert Log and transmission of alerts (traps) to the management server (in cases where a network management tool such as ProCurve Manager is in use). Use this option when you don't want to use the Alert Log.

The Fault Detection Window also contains three Change Control Buttons:

- **Apply Changes.** This button stores the settings you have selected for all future sessions with the web browser interface until you decide to change them.
- **Clear Changes.** This button removes your settings and returns the settings for the list box to the level it was at in the last saved detection-setting session.
- **Reset to Default Settings.** This button reverts the policy setting to Medium Sensitivity for Log Network Problems.

Using the ProCurve Web Browser Interface
Status Reporting Features

Switch Memory and Configuration

Contents

Overview	6-3
Configuration File Management	6-3
Using the CLI To Implement Configuration Changes	6-6
Using the Menu and Web Browser Interfaces To Implement Configuration Changes	6-10
Menu: Implementing Configuration Changes	6-10
Using Save and Cancel in the Menu Interface	6-10
Rebooting from the Menu Interface	6-11
Web: Implementing Configuration Changes	6-13
Using Primary and Secondary Flash Image Options	6-14
Displaying the Current Flash Image Data	6-14
Switch Software Downloads	6-16
Local Switch Software Replacement and Removal	6-17
Rebooting the Switch	6-19
Operating Notes about Booting	6-19
Boot and Reload Command Comparison	6-20
Setting the Default Flash	6-21
Booting from the Default Flash (Primary or Secondary)	6-22
Booting from a Specified Flash	6-23
Using Reload	6-24
Multiple Configuration Files	6-26
General Operation	6-27
Transitioning to Multiple Configuration Files	6-29
Listing and Displaying Startup-Config Files	6-30
Viewing the Startup-Config File Status with Multiple Configuration Enabled	6-30
Displaying the Content of A Specific Startup-Config File	6-31

Changing or Overriding the Reboot Configuration Policy	6-31
Managing Startup-Config Files in the Switch	6-33
Renaming an Existing Startup-Config File	6-34
Creating a New Startup-Config File	6-34
Erasing a Startup-Config File	6-35
Using the Clear + Reset Button Combination To Reset the Switch to Its Default Configuration	6-37
Transferring Startup-Config Files To or From a Remote Server	6-38
TFTP: Copying a Configuration File to a Remote Host	6-38
TFTP: Copying a Configuration File from a Remote Host	6-39
Xmodem: Copying a Configuration File to a Serially Connected Host	6-40
Xmodem: Copying a Configuration from a Serially Connected Host	6-40
Operating Notes for Multiple Configuration Files	6-40
Automatic Configuration Update with DHCP Option 66	6-41
CLI Command	6-41
Possible Scenarios for Updating the Configuration File	6-42
Operating Notes	6-42
Log Messages	6-43

Overview

This chapter describes:

- How switch memory manages configuration changes
 - How the CLI implements configuration changes
 - How the menu interface and web browser interface implement configuration changes
 - How the switch provides software options through primary/secondary flash images
 - How to use the switch's primary and secondary flash options, including displaying flash information, booting or restarting the switch, and other topics
-

Configuration File Management

The switch maintains two configuration files, the *running-config* file and the *startup-config* file.




Figure 6-1. Conceptual Illustration of Switch Memory Operation

- **Running Config File:** Exists in volatile memory and controls switch operation. If no configuration changes have been made in the CLI since the switch was last booted, the running-config file is identical to the startup-config file.
-

- **Startup-config File:** Exists in flash (non-volatile) memory and is used to preserve the most recently-saved configuration as the “permanent” configuration.

Booting the switch replaces the current running-config file with a new running-config file that is an exact copy of the current startup-config file.

Note

Any of the following actions boots the switch:

- Executing the **boot** or the **reload** command in the CLI
- Executing the **boot** command in the menu interface
- Pressing the Reset button on the front of the switch
- Removing, then restoring power to the switch

For more on reboots and the switch’s dual-flash images, refer to “Using Primary and Secondary Flash Image Options” on page 6-14.

Options for Saving a New Configuration. Making one or more changes to the running-config file creates a new operating configuration. *Saving* a new configuration means to overwrite (replace) the current startup-config file with the current running-config file. This means that if the switch subsequently reboots for any reason, it will resume operation using the new configuration instead of the configuration previously defined in the startup-config file. There are three ways to save a new configuration:

- **In the CLI:** Use the **write memory** command. This overwrites the current startup-config file with the contents of the current running-config file.
- **In the menu interface:** Use the **Save** command. This overwrites *both* the running-config file and the startup-config file with the changes you have specified in the menu interface screen.
- **In the web browser interface:** Use the **[Apply Changes]** button or other appropriate button. This overwrites *both* the running-config file and the startup-config file with the changes you have specified in the web browser interface window.

Note that using the CLI instead of the menu or web browser interface gives you the option of changing the running configuration without affecting the startup configuration. This allows you to test the change without making it “permanent”. When you are satisfied that the change is satisfactory, you can make it permanent by executing the **write memory** command. For example, suppose you use the following command to disable port 5:

```
ProCurve(config)# interface ethernet 5 disable
```

The above command disables port 5 in the running-config file, but not in the startup-config file. Port 5 remains disabled only until the switch reboots. If you want port 5 to remain disabled through the next reboot, use **write memory** to save the current running-config file to the startup-config file in flash memory.

```
ProCurve(config)# write memory
```

If you use the CLI to make a configuration change and then change from the CLI to the Menu interface without first using write memory to save the change to the startup-config file, then the switch prompts you to save the change. For example, if you use the CLI to create VLAN 20, and then select the menu interface, VLAN 20 is configured in the running-config file, but not in the startup-config file. In this case you will see:

```
ProCurve(config)# vlan 20
ProCurve(config)# menu
Do you want to save current configuration [y/n]?
```

If you type **[Y]**, the switch overwrites the startup-config file with the running-config file, and your configuration change(s) will be preserved across reboots. If you type **[N]**, your configuration change(s) will remain only in the running-config file. In this case, if you do not subsequently save the running-config file, your unsaved configuration changes will be lost if the switch reboots for any reason.

Storing and Retrieving Configuration Files. You can store or retrieve a backup copy of the startup-config file on another device. For more information, refer to the section on “Transferring Switch Configurations” on page A-29 in Appendix A on “File Transfers”.

USB Autorun. This feature supports the ability to auto execute CLI commands stored on a USB flash drive (for example, to configure the switch, update software, retrieve diagnostics, etc.). For more information, refer to the section on “Using USB Autorun” on page A-47.

Using the CLI To Implement Configuration Changes

The CLI offers these capabilities:

- Access to the full set of switch configuration features
- The option of testing configuration changes before making them permanent

How To Use the CLI To View the Current Configuration Files. Use **show** commands to view the configuration for individual features, such as port status or Spanning Tree Protocol. However, to view either the entire startup-config file or the entire running-config file, use the following commands:

- **show config** — Displays a listing of the current startup-config file.
- **show running-config** — Displays a listing of the current running-config file.
- **write terminal** — Displays a listing of the current running-config file.
- **show config status** — Compares the startup-config file to the running-config file and lists one of the following results:
 - If the two configurations are the same you will see:
 - Running configuration is the same as the startup configuration.
 - If the two configurations are different, you will see:
 - Running configuration has been changed and needs to be saved.

Note

Show config, **show running-config**, and **write terminal** commands display the configuration settings that differ from the switch's factory-default configuration.

How To Use the CLI To Reconfigure Switch Features. Use this procedure to permanently change the switch configuration (that is, to enter a change in the startup-config file).

1. Use the appropriate CLI commands to reconfigure the desired switch parameters. This updates the selected parameters in the running-config file.
2. Use the appropriate **show** commands to verify that you have correctly made the desired changes.

3. Observe the switch's performance with the new parameter settings to verify the effect of your changes.
4. When you are satisfied that you have the correct parameter settings, use the **write memory** command to copy the changes to the startup-config file.

Syntax: write memory

Saves the running configuration file to the startup-config. The saved configuration becomes the boot-up configuration of the switch on the next boot.

When using redundant management, saves the running configuration of the switch to flash on the active management module. The saved configuration becomes the boot-up configuration of the switch the next time it is booted. The saved configuration file is sync'd to the standby management module.

Note: If the active management module and the standby management module are running on different operating systems because the boot set-default command was executed and then the standby module was rebooted, the write memory command displays this warning: "Warning: The next reboot or failover is set to boot from a different software image. These config changes may be incompatible or not used after a reboot or failover."

For example, the default port mode setting is **auto**. Suppose that your network uses Cat 3 wiring and you want to connect the switch to another autosensing device capable of 100 Mbps operation. Because 100 Mbps over Cat 3 wiring can introduce transmission problems, the recommended port mode is **auto-10**, which allows the port to negotiate full- or half-duplex, but restricts speed to 10 Mbps. The following command configures port A5 to auto-10 mode in the running-config file, allowing you to observe performance on the link without making the mode change permanent.

```
ProCurve(config)# interface e a5 speed-duplex auto-10
```

After you are satisfied that the link is operating properly, you can save the change to the switch's permanent configuration (the startup-config file) by executing the following command:

```
ProCurve(config)# write memory
```

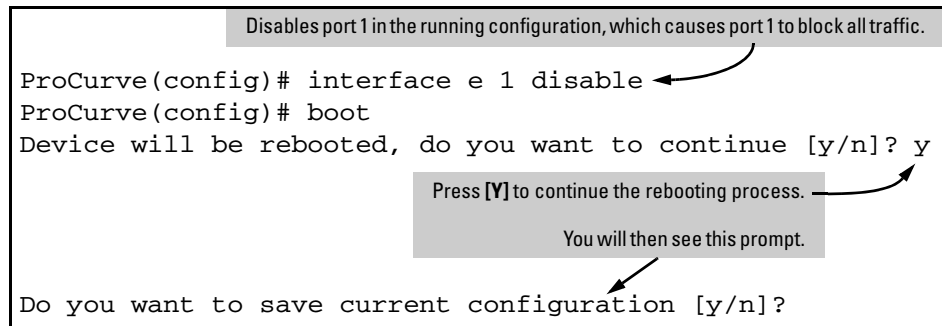
The new mode (**auto-10**) on port A5 is now saved in the startup-config file, and the startup-config and running-config files are identical. If you subsequently reboot the switch, the **auto-10** mode configuration on port A5 will remain because it is included in the startup-config file.

How To Cancel Changes You Have Made to the Running-Config File.

If you use the CLI to change parameter settings in the running-config file, and then decide that you don't want those changes to remain, you can use either of the following methods to remove them:

- Manually enter the earlier values you had for the changed settings. (This is recommended if you want to restore a small number of parameter settings to their previous boot-up values.)
- Update the running-config file to match the startup-config file by rebooting the switch. (This is recommended if you want to restore a larger number of parameter settings to their previous boot-up values.)

If you use the CLI to change a parameter setting, and then execute the **boot** command without first executing the **write memory** command to save the change, the switch prompts you to specify whether to save the changes in the current running-config file. For example:



```
ProCurve(config)# interface e 1 disable
ProCurve(config)# boot
Device will be rebooted, do you want to continue [y/n]? y
Do you want to save current configuration [y/n]?
```

Disables port 1 in the running configuration, which causes port 1 to block all traffic.

Press [Y] to continue the rebooting process.

You will then see this prompt.

Figure 6-2. Boot Prompt for an Unsaved Configuration

The above prompt means that one or more parameter settings in the running-config file differ from their counterparts in the startup-config file and you need to choose which config file to retain and which to discard.

- If you want to update the startup-config file to match the running-config file, press **[Y]** for “yes”. (This means that the changes you entered in the running-config file will be saved in the startup-config file.)
- If you want to discard the changes you made to the running-config file so that it will match the startup-config file, then press **[N]** for “no”. (This means that the switch will discard the changes you entered in the running-config file and will update the running-config file to match the startup-config file.)

Note

If you use the CLI to make a change to the running-config file, you should either use the **write memory** command or select the save option allowed during a reboot (figure 6-6-2, above) to save the change to the startup-config file. That is, if you use the CLI to change a parameter setting, but then reboot the switch from either the CLI or the menu interface without first executing the **write memory** command in the CLI, the current startup-config file will replace the running-config file, and any changes in the running-config file will be lost.

Using the **Save** command in the menu interface does not save a change made to the running config by the CLI unless you have also made a configuration change in the menu interface. Also, the menu interface displays the current running-config values. Thus, where a parameter setting is accessible from both the CLI and the menu interface, if you change the setting in the CLI, the new value will appear in the menu interface display for that parameter. *However, as indicated above, unless you also make a configuration change in the menu interface, only the write memory command in the CLI will actually save the change to the startup-config file.*

How To Reset the startup-config and running-config Files to the Factory Default Configuration. This command reboots the switch, replacing the contents of the current startup-config and running-config files with the factory-default startup configuration.

Syntax: erase startup-config

For example:

```
ProCurve(config)# erase startup-config
Configuration will be deleted and device rebooted, continue [y/n]?
```

Figure 6-3. Example of erase startup-config Command

Press [y] to replace the current configuration with the factory default configuration and reboot the switch. Press [n] to retain the current configuration and prevent a reboot.

In a redundant management system, this command erases the startup config file on both the active and the standby management modules as long as redundancy has not been disabled. If the standby management module is not in standby mode or has failed selftest, the startup config file is not erased.

Using the Menu and Web Browser Interfaces To Implement Configuration Changes

The menu and web browser interfaces offer these advantages:

- Quick, easy menu or window access to a subset of switch configuration features
- Viewing several related configuration parameters in the same screen, with their default and current settings
- Immediately changing both the running-config file and the startup-config file with a single command

Menu: Implementing Configuration Changes

You can use the menu interface to simultaneously save and implement a subset of switch configuration changes without having to reboot the switch. That is, when you save a configuration change in the menu interface, you simultaneously change both the running-config file and the startup-config file.

Note

The only exception to this operation are two VLAN-related parameter changes that require a reboot—described under “Rebooting To Activate Configuration Changes” on page 6-12.

Using **S**ave and **C**ancel in the Menu Interface

For any configuration screen in the menu interface, the Save command:

1. Implements the changes in the running-config file
2. Saves your changes to the startup-config file

If you decide not to save and implement the changes in the screen, select **C**ancel to discard them and continue switch operation with the current operation. For example, suppose you have made the changes shown below in the System Information screen:

To save and implement the changes for all parameters in this screen, press the **[Enter]** key, then press **[S]** (for **Save**). To cancel all changes, press the **[Enter]** key, then press **[C]** (for **Cancel**)

```

ProCurve
----- CONSOLE - MANAGER MODE -----
                Switch Configuration - System Information

System Name :    ProCurve Switch
System Contact :
System Location :

Inactivity Timeout (min) [0] : 0      MAC Age Time (sec) [300] : 300
Inbound Telnet Enabled [Yes] : Yes    Web Agent Enabled [Yes] : Yes
Time Sync Method [None] : TIMEP
TimeP Mode [Disabled] : Disabled

Time Zone [0] : 0
Daylight Time Rule [None] : Continental-US-and-Canada

Actions->  Cancell  Edit   Save   Help

Select Daylight Time Rule for your location.
Use arrow keys to change field selection, <Space> to toggle field choices,
and <Enter> to go to Actions.

```

Figure 6-4. Example of Pending Configuration Changes You Can Save or Cancel

Note

If you reconfigure a parameter in the CLI and then go to the menu interface without executing a **write memory** command, those changes are stored only in the running configuration (even if you execute a Save operation in the menu interface). If you then execute a switch **boot** command in the menu interface, the switch discards the configuration changes made while using the CLI. To ensure that changes made while using the CLI are saved, execute **write memory** in the CLI before rebooting the switch.

Rebooting from the Menu Interface

- Terminates the current session and performs a reset of the operating system
- Activates any configuration changes that require a reboot
- Resets statistical counters to zero

(Note that statistical counters can be reset to zero without rebooting the switch. See “To Display the Port Counter Summary Report” on page 18.)

To Reboot the switch, use the **Reboot Switch** option in the Main Menu. (Note that the Reboot Switch option is not available if you log on in Operator mode; that is, if you enter an Operator password instead of a manager password at the password prompt.)

Switch Memory and Configuration

Using the Menu and Web Browser Interfaces To Implement Configuration Changes

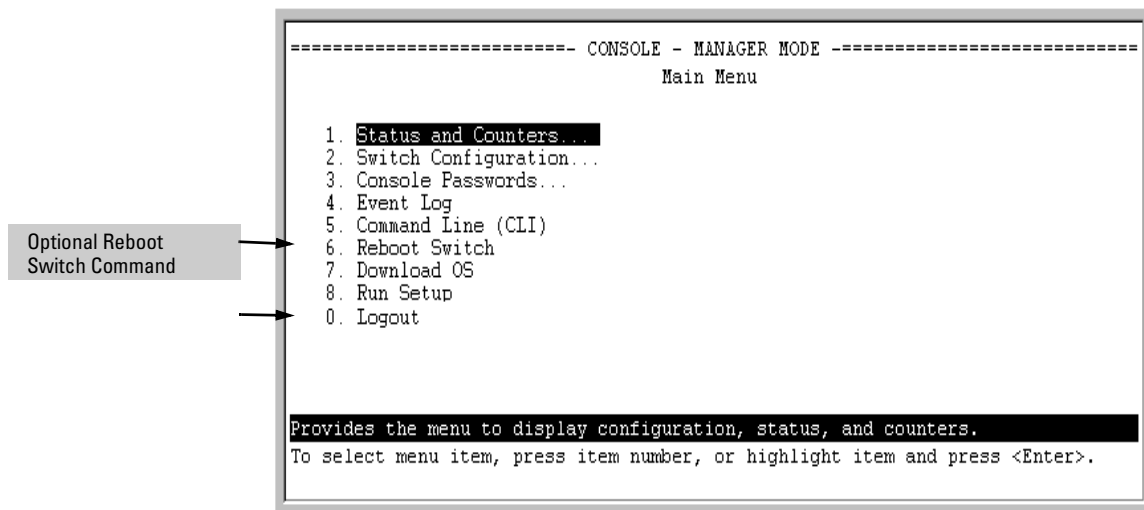


Figure 6-5. The Reboot Switch Option in the Main Menu

Rebooting To Activate Configuration Changes. Configuration changes for most parameters become effective as soon as you save them. However, you must reboot the switch in order to implement a change in the **Maximum VLANs to support** parameter.

(To access these parameters, go to the Main menu and select **2. Switch Configuration**, then **8. VLAN Menu**, then **1. VLAN Support**.)

If configuration changes requiring a reboot have been made, the switch displays an asterisk (*) next to the menu item in which the change has been made. For example, if you change and save parameter values for the **Maximum VLANs to support** parameter, an asterisk appears next to the **VLAN Support** entry in the VLAN Menu screen, and also next to the **Switch Configuration** ..entry in the Main menu, as shown in Figure 6-6:

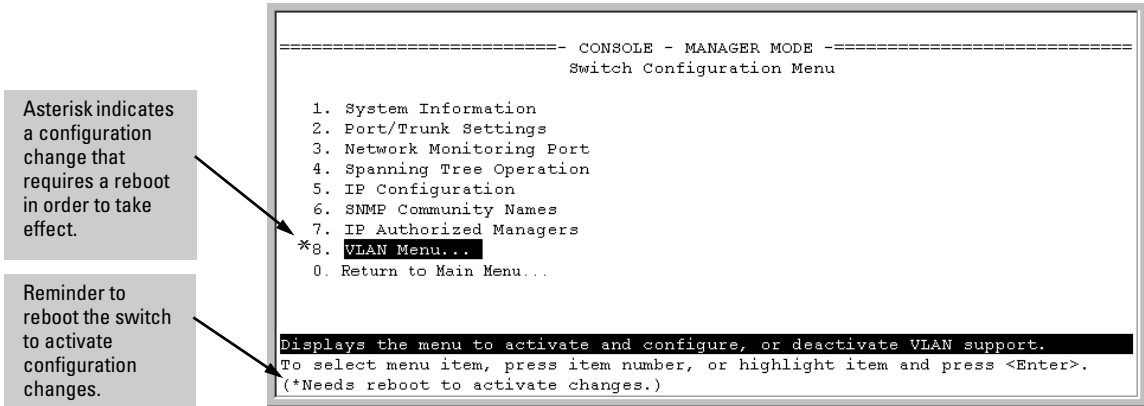


Figure 6-6. Indication of a Configuration Change Requiring a Reboot

Web: Implementing Configuration Changes

You can use the web browser interface to simultaneously save and implement a subset of switch configuration changes without having to reboot the switch. That is, when you save a configuration change (in most cases, by clicking on **[Apply Changes]** or **[Apply Settings]**), you simultaneously change both the running-config file and the startup-config file.

Note

If you reconfigure a parameter in the CLI and then go to the browser interface without executing a **write memory** command, those changes will be saved to the startup-config file if you click on **[Apply Changes]** or **[Apply Settings]** in the web browser interface.

Using Primary and Secondary Flash Image Options

The switches covered in this guide feature two flash memory locations for storing switch software image files:

- **Primary Flash:** The default storage for a switch software image.
- **Secondary Flash:** The additional storage for either a redundant or an alternate switch software image.

With the Primary/Secondary flash option you can test a new image in your system without having to replace a previously existing image. You can also use the image options for troubleshooting. For example, you can copy a problem image into Secondary flash for later analysis and place another, proven image in Primary flash to run your system. The switch can use only one image at a time.

The following tasks involve primary/secondary flash options:

- Displaying the current flash image data and determining which switch software versions are available
- Switch software downloads
- Replacing and removing (erasing) a local switch software version
- System booting

Displaying the Current Flash Image Data

Use the commands in this section to:

- Determine whether there are flash images in both primary and secondary flash
- Determine whether the images in primary and secondary flash are the same
- Identify which switch software version is currently running

Viewing the Currently Active Flash Image Version. This command identifies the software version on which the switch is currently running, and whether the active version was booted from the primary or secondary flash image.

Syntax:show version

For example, if the switch is using a software version of K.12.XX stored in Primary flash, **show version** produces the following:

```
ProCurve(config)# show version

Image stamp:   /su/code/build/info(s01)
               Dec 01 2006 10:50:26
               K.12.XX
               1223
Boot Image:    Primary
```

Figure 6-7. Example Showing the Identity of the Current Flash Image

Determining Whether the Flash Images Are Different Versions. If the flash image sizes in primary and secondary are the same, then in almost every case, the primary and secondary images are identical. This command provides a comparison of flash image sizes, plus the boot ROM version and from which flash image the switch booted. For example, in the following case, the images are different versions of the switch software, and the switch is running on the version stored in the secondary flash image:

```
ProCurve(config)# show flash
Image          Size(Bytes)   Date   Version  Build #
-----
Primary Image  : 7493854    03/21/07 K.12.29  1617
Secondary Image : 7463821    03/23/07 K.12.30  1700

Boot Rom Version: K.12.30
Default Boot    : Primary
```

Will boot from primary flash
on the next boot.

Figure 6-8. Example Showing Different Flash Image Versions

Determining Which Flash Image Versions Are Installed. The **show version** command displays which software version the switch is currently running and whether that version booted from primary or secondary flash. Thus, if the switch booted from primary flash, you will see the version number of the software version stored in primary flash, and if the switch booted from secondary flash, you will see the version number of the software version stored in secondary flash. Thus, by using **show version**, then rebooting the switch from the opposite flash image and using **show version** again, you can determine the version(s) of switch software in both flash sources. For example:

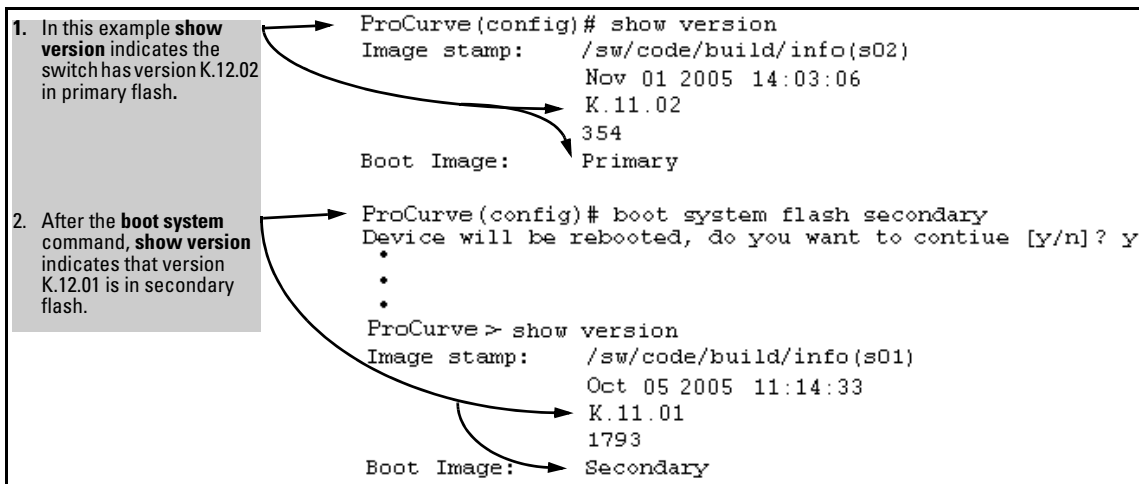


Figure 6-9. Determining the Software Version in Primary and Secondary Flash

Switch Software Downloads

The following table shows the switch’s options for downloading a software version to flash and booting the switch from flash

Table 6-1. Primary/Secondary Memory Access

Action	Menu	CLI	Web Browser	SNMP
Download to Primary	Yes	Yes	Yes	Yes
Download to Secondary	No	Yes	No	Yes
Boot from Primary	Yes	Yes	Yes	Yes
Boot from Secondary	No	Yes	No	Yes

The different software download options involve different **copy** commands, plus **xmodem**, **usb**, and **ftpp**. These topics are covered in Appendix A, “File Transfers”.

Download Interruptions. In most cases, if a power failure or other cause interrupts a flash image download, the switch reboots with the image previously stored in primary flash. In the unlikely event that the primary image is corrupted, as a result of an interruption, the switch will reboot from secondary flash and you can either copy the secondary image into primary or download another image to primary from an external source. Refer to Appendix A, “File Transfers”.

Local Switch Software Replacement and Removal

This section describes commands for erasing a software version and copying an existing software version between primary and secondary flash.

Note

It is not necessary to erase the content of a flash location before downloading another software file. The process automatically overwrites the previous file with the new file. If you want to remove an unwanted software version from flash, ProCurve recommends that you do so by overwriting it with the same software version that you are using to operate the switch, or with another acceptable software version. To copy a software file between the primary and secondary flash locations, refer to “Copying a Switch Software Image from One Flash Location to Another”, below.

The local commands described here are for flash image management within the switch. To download a software image file from an external source, refer to Appendix A, “File Transfers”.

Copying a Switch Software Image from One Flash Location to

Another. When you copy the flash image from primary to secondary or the reverse, the switch overwrites the file in the destination location with a copy of the file from the source location. This means you *do not* have to erase the current image at the destination location before copying in a new image.

Caution

Verify that there is an acceptable software version in the source flash location from which you are going to copy. Use the **show flash** command or, if necessary, the procedure under “Determining Which Flash Image Versions Are Installed” on page 6-15 to verify an acceptable software version. Attempting to copy from a source image location that has a corrupted flash image overwrites the image in the destination flash location. In this case, the switch will not have a valid flash image in either flash location, but will continue running on a temporary flash image in RAM. *Do not reboot the switch.* Instead, immediately download another valid flash image to primary or secondary flash. Otherwise, if the switch is rebooted without a software image in either primary or secondary flash, the temporary flash image in RAM will be cleared and the switch will go down. To recover, refer to “Restoring a Flash Image” on page C-84 (in the “Troubleshooting” Appendix).

Syntax: copy flash flash <destination flash>

where: *destination flash* = **primary** or **secondary**:

For example, to copy the image in secondary flash to primary flash:

1. Verify that there is a valid flash image in the secondary flash location. The following figure indicates that a software image is present in secondary flash. (If you are unsure whether the image in secondary flash is valid, try booting from it before you proceed, by using **boot system flash secondary**.)

```
ProCurve(config)# show flash
Image          Size(Bytes)  Date      Version
-----
Primary Image  : 3275389   11/05/05  K.11.30
Secondary Image: 3258128   10/25/05  K.11.20
Boot Rom Version: K.11.Z3
Current Boot   : Primary
```

The unequal code size, differing dates, and differing version numbers indicates two different versions of the software.

Figure 6-10. Example Indicating Two Different Software Versions in Primary and Secondary Flash

Execute the copy command as follows:

```
ProCurve(config)# copy flash flash primary
```

Erasing the Contents of Primary or Secondary Flash. This command deletes the software image file from the specified flash location.

Caution:

No Undo!

Before using this command in one flash image location (primary or secondary), ensure that you have a valid software file in the other flash image location (secondary or primary). If the switch has only one flash image loaded (in either primary or secondary flash) and you erase that image, then the switch does not have a software image stored in flash. In this case, if you do not reboot or power cycle the switch, you can recover by using xmodem or tftp to download another software image.

Syntax: erase flash < primary | secondary >

For example, to erase the software image in primary flash, do the following:

1. First verify that a usable flash image exists in secondary flash. The most reliable way to ensure this is to reboot the switch from the flash image you want to retain. For example, if you are planning to erase the primary image, then first reboot from the secondary image to verify that the secondary image is present and acceptable for your system:

```
ProCurve# boot system flash secondary
```

2. Then erase the software image in the selected flash (in this case, primary):

```
ProCurve# erase flash primary
The Primary OS Image will be deleted, continue [y/n]? _
```

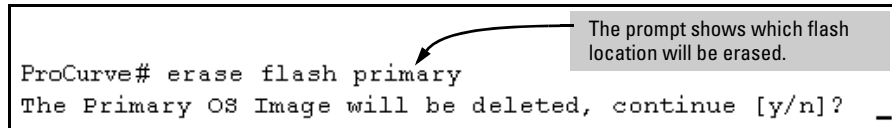


Figure 6-11. Example of Erase Flash Prompt

3. Type **y** at the prompt to complete the flash erase.
4. Use **show flash** to verify erasure of the selected software flash image

```
ProCurve# show flash
Compressed Primary Code size    = 0
Compressed Secondary Code size  = 2555802
Boot Rom Version:               E.05.04
Current Boot:                   Secondary
```

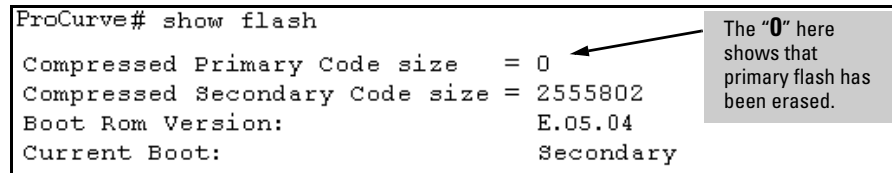


Figure 6-12. Example of Show Flash Listing After Erasing Primary Flash

In redundant management systems, this command will erase the selected flash in both the active and the standby management modules. If redundancy has been disabled or the standby module has failed selftest, this command only affects the active management module.

Rebooting the Switch

Operating Notes about Booting

Default Boot Source. The switch reboots from primary flash by default unless you specify the secondary flash by entering either the **boot system flash [primary | secondary]** or **boot set-default flash [primary | secondary]** command. Both the **boot** command and the **reload** command will reboot based on how these options have been selected.

Boot Attempts from an Empty Flash Location. In this case, the switch aborts the attempt and displays

```
Image does not exist
Operation aborted.
```

Interaction of Primary and Secondary Flash Images with the Current Configuration. The switch has one startup-config file (page 6-3), which it always uses for reboots, regardless of whether the reboot is from primary or secondary flash. Also, for rebooting purposes, it is not necessary for the software image and the startup-config file to support identical software fea-

tures. For example, suppose you have just downloaded a software upgrade that includes new features that are not supported in the software you used to create the current startup-config file. In this case, the software simply assigns factory-default values to the parameters controlling the new features. Similarly, If you create a startup-config file while using a version “Y” of the switch software, and then reboot the switch with an earlier software version “X” that does not include all of the features found in “Y”, the software simply ignores the parameters for any features that it does not support.

Scheduled Reload. If no parameters are entered after the **reload** command, an immediate reboot is executed. The **reload at** and **reload after** command information is not saved across reboots. If the switch is rebooted before a scheduled reload command is executed, the command is effectively cancelled. When entering a **reload at** or **reload after** command, a prompt will appear to confirm the command before it can be processed by the switch. For the **reload at** command, if *mm/dd/yy* are left blank, the current day is assumed.

The scheduled reload feature removes the requirement to physically reboot the switch at inconvenient times (for example, at 1:00 in the morning). Instead, a **reload at 1:00 mm/dd** command can be executed (where *mm/dd* is the date the switch is scheduled to reboot).

Boot and Reload Command Comparison

The switch offers reboot options through the **boot** and **reload** commands, plus the options inherent in a dual-flash image system. Generally, using **boot** provides more comprehensive self-testing; using **reload** gives you a faster reboot time.

Table 6-2. Comparing the Boot and Reload Commands

Actions	Included In Boot?	Included In Reload	Note
Save all configuration changes since the last boot or reload	Optional, with prompt	Optional with reload <cr>, when prompt displays. Not saved with reload at/after commands; No prompt is displayed.	Config changes saved to the startup-config file if "y" is selected (reload command).
Perform all system self-tests	Yes	No	The reload command provides a faster system reboot.
Choice of primary or secondary flash image	Yes	No—Uses the current flash image.	
Perform a scheduled reboot	No	Yes	Use the reload command with after/at parameters (see page 6-25 for details).

Setting the Default Flash

You can specify the default flash to boot from on the next boot by entering the **boot set-default flash** command.

Syntax: boot set-default flash [primary |secondary]

Upon booting, set the default flash for the next boot to primary or secondary.

```
ProCurve(config)# boot set-default flash secondary
ProCurve(config)# show flash
Image           Size(Bytes)   Date    Version  Build #
-----
Primary Image   : 7476770    03/15/07 K.12.XX   64
Secondary Image : 7476770    03/15/07 K.12.XX   64
Boot Rom Version: K.12.02
Default Boot    : Secondary

ProCurve(config)# boot
This management module will now reboot from secondary and will become
the standby module! You will need to use the other management module's
console interface. Do you want to continue [y/n]?
```

Figure 6-13. Example of boot set-default Command with Default Flash Set to Secondary (with a Redundant Management Module Present)

Booting from the Default Flash (Primary or Secondary)

The **boot** command boots the switch from the flash image that you are currently booted on, or the flash image that was set either by the **boot set-default** command or by the last executed **boot system flash <primary | secondary>** command. This command also executes the complete set of subsystem self-tests. You have the option of specifying a configuration file.

Syntax: boot [system [flash <primary | secondary>]] [config FILENAME]

Reboots the switch from the flash that you are currently booted on (primary or secondary). You can select which image to boot from during the boot process itself. When using redundant management, the switch will failover to the standby management module.

Note: *This is changed from always booting from primary flash. You are prompted with a message which will indicate the flash being booted from.*

system: *Boots the switch. You can specify the flash image to boot from. When using redundant management, boots both the active and standby management modules.*

config: *You can optionally select a configuration file from which to boot.*

```
ProCurve(config)# boot
This management module will now reboot from primary image and will become
the standby module! You will need to use the other management module's
console interface. Do you want to continue [y/n]? y

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

Figure 6-14. Example of Boot Command (Default Primary Flash) with Redundant Management

In the above example, typing either a **y** or **n** at the second prompt initiates the reboot operation. (Entering **y** saves any configuration changes from the running-config file to the startup-config file; entering **n** discards them.)

```
ProCurve(config)# show flash
Image           Size(Bytes)   Date   Version  Build #
-----
Primary Image   : 7497114   03/29/07 K.12.XX   57
Secondary Image : 7497114   03/29/07 K.12.XX   57
Boot Rom Version: K.12.03
Default Boot    : Primary

ProCurve(config)# boot set-default flash secondary
This command changes the location of the default boot. This command will
change the default flash image to boot from secondary. Hereafter, 'reload'
'boot' commands will boot from secondary. Do you want to continue [y/n]? y

ProCurve(config)# boot
This management module will now reboot from secondary image and will become
the standby module! You will need to use the other management module's
console interface. Do you want to continue [y/n]? n
```

Figure 6-15. Example of Boot Command Booting from a Different Flash than the Current Flash (with Redundant Management Module Present)

Booting from a Specified Flash

This version of the boot command gives you the option of specifying whether to reboot from primary or secondary flash, and is the required command for rebooting from secondary flash. This option also executes the complete set of subsystem self-tests.

Syntax: boot system flash < primary | secondary >

For example, to reboot the switch from secondary flash when there are no pending configuration changes in the running-config file:

```
ProCurve(config)# boot system flash secondary
System will be rebooted from secondary image. Do you want to continue [y/n]?
```

Figure 6-16. Example of Boot Command with Secondary Flash Option

In the above example, typing either a **y** or **n** at the second prompt initiates the reboot operation.

Using the Fastboot feature. The **fastboot** command allows a boot sequence that skips the internal power-on self-tests, resulting in a faster boot time. When using redundant management and fastboot is enabled, it is saved to the standby management module when the config files are synchronized. Fastboot is used during the next bootup on either management module.

Syntax: [no] fastboot

Enables the fastboot option

*The **no** option disables the feature.*

Syntax: show fastboot

Shows the status of the fastboot feature, either enabled or disabled.

The fastboot command is shown below.

```
ProCurve(config)# fastboot
```

Using Reload

The **Reload** command reboots the switch from the flash image that you are currently booted on (primary or secondary) or the flash image that was set either by the **boot set-default** command or by the last executed **boot system flash <primary | secondary>** command. Because **reload** bypasses some subsystem self-tests, the switch reboots faster than if you use either of the **boot** command options. If you are using redundant management and redundancy is enabled, the switch will failover to the other management module.

Syntax: reload

For example, if you change the number of VLANs the switch supports, you must reboot the switch in order to implement the change. The **reload** command prompts you to save or discard the configuration changes.

```
ProCurve(config)# max-vlans 12
Command will take effect after saving configuration and reboot.

ProCurve(config)# reload
This command will cause a switchover to the other management module
which may not be running the same software image and configurations.
Do you want to continue [y/n]? y
```

Figure 6-17. Using Reload with Redundant Management and Pending Configuration Changes

Scheduled Reload. Beginning with software release K.11.34, additional parameters have been added to the **reload** command to allow for a scheduled reboot of the switch via the CLI.

Syntax: [no] reload [after <[dd:]hh:]mm> | at <hh:mm[:ss]> [<mm/dd/[yy]yy>]]

Enables a scheduled warm reboot of the switch. The switch boots up with the same startup config file and using the same flash image as before the reload.

Caution: *When using redundant management, the **reload at/after** command causes a switchover at the scheduled time to the other management module, which may not be running the same software image or have the same configurations.*

Parameters include:

- **after:** *Schedules a warm reboot of the switch after a given amount of time has passed.*
- **at:** *Schedules a warm reboot of the switch at a given time.*

*The **no** form of the command removes a pending reboot request.*

For more details and examples, see below.

The scheduled reload feature removes the requirement to physically reboot the switch at inconvenient times (for example, at 1:00 in the morning). Instead, a **reload at 1:00 mm/dd** command can be executed (where *mm/dd* is the date the switch is scheduled to reboot).

Note

Configuration changes are not saved with **reload at** or **reload after** commands. No prompt to save configuration file changes is displayed. See Table 6-2 on page 6-21.

Examples of scheduled **reload** commands:

- To schedule a reload in 15 minutes:
ProCurve# reload after 15
- To schedule a reload in 3 hours:
ProCurve# reload after 03:00
- To schedule a reload for the same time the following day:
ProCurve# reload after 01:00:00
- To schedule a reload for the same day at 12:05:
ProCurve# reload at 12:05
- To schedule a reload on some future date:
ProCurve# reload at 12:05 01/01/2008

```
ProCurve(config)# reload after 04:14:00
Reload scheduled in 4 days, 14 hours, 0 minutes
This command will cause a switchover at the scheduled time to the
other management module which may not be running the same software
image and configurations. Do you want to continue [y/n]?
```

Figure 6-18. An Example of the reload Command with a Redundant Management System

Multiple Configuration Files

Action	Page
Listing and Displaying Startup-Config Files	6-30
Changing or Overriding the Reboot Configuration Policy	6-31
Managing Startup-Config Files	
Renaming Startup-Config Files	6-34
Copying Startup-Config Files	6-34
Erasing Startup-Config Files	6-35
Effect of Using the Clear + Reset Buttons	6-37
Copying Startup-Config Files to or from a Remote Server	6-38

This method of operation means that you cannot preserve different startup-config files across a reboot without using remote storage.

The switch allows up to three startup-config files with options for selecting which startup-config file to use for:

- A fixed reboot policy using a specific startup-config file for a specific boot path (primary or secondary flash)
- Overriding the current reboot policy on a per-instance basis




Figure 6-19. Optional Reboot Process

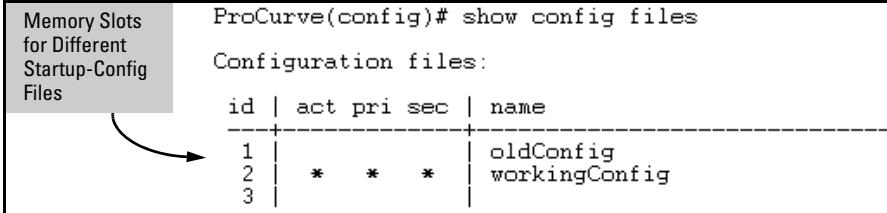
While you can still use remote storage for startup-config files, you can now maintain multiple startup-config files on the switch and choose which version to use for a reboot policy or an individual reboot.

This choice of which configuration file to use for the startup-config at reboot provides the following new options:

- The switch can reboot with different configuration options without having to exchange one configuration file for another from a remote storage location.
- Transitions from one software release to another can be performed while maintaining a separate configuration for the different software release versions.
- By setting a reboot policy using a known good configuration and then overriding the policy on a per-instance basis, you can test a new configuration with the provision that if an unattended reboot occurs, the switch will come up with the known, good configuration instead of repeating a reboot with a misconfiguration.

General Operation

Multiple Configuration Storage in the Switch. The switch uses three memory “slots”, with identity (**id**) numbers of **1**, **2**, and **3**.



```
ProCurve(config)# show config files
Configuration files:
  id | act pri sec | name
-----|-----|-----
  1 |          | oldConfig
  2 | * * *      | workingConfig
  3 |          |
```

The screenshot shows a terminal window with the command 'ProCurve(config)# show config files' and its output. The output is a table with columns 'id', 'act', 'pri', 'sec', and 'name'. Row 1 has '1' in the 'id' column and 'oldConfig' in the 'name' column. Row 2 has '2' in the 'id' column and '*' in the 'act', 'pri', and 'sec' columns, with 'workingConfig' in the 'name' column. Row 3 has '3' in the 'id' column. A callout box on the left, titled 'Memory Slots for Different Startup-Config Files', has an arrow pointing to the 'id' column.

A startup-config file stored in a memory slot has a unique, changeable file name. The switches covered in this guide can use the startup-config in any of the memory slots (if the software version supports the configured features).

Boot Options. With multiple startup-config files in the switch you can specify a policy for the switch to use upon reboot. The options include:

- Use the designated startup-config file with either or both reboot paths (primary or secondary flash)
- Override the current reboot policy for one reboot instance by specifying a boot path (primary or secondary flash) and the startup-config file to use.

Changing the Startup-Config File. When the switch reboots, the startup-config file supplies the configuration for the running-config file the switch uses to operate. Making changes to the running-config file and then executing a **write-mem** command (or, in the Menu interface, the **Save** command) are written back to the startup-config file used at the last reboot. For example, suppose that a system administrator performs the following on a switch that has two startup-config files (**workingConfig** and **backupConfig**):

1. Reboot the switch through the Primary boot path using the startup-config file named **backupConfig**.
2. Use the CLI to make configuration changes in the running-config file, and then execute **write mem**.

The result is that the startup-config file used to reboot the switch is modified by the actions in step 2.

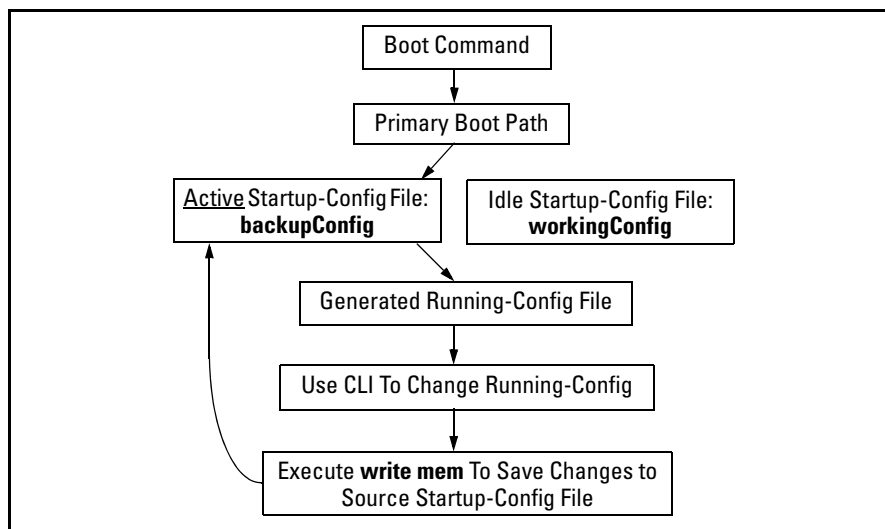


Figure 6-20. Example of Reboot Process and Making Changes to the Startup-Config File

Creating an Alternate Startup-Config File. There are two methods for creating a new configuration file:

- Copy an existing startup-config file to a new filename, then reboot the switch, make the desired changes to the running-config file, then execute **write memory**. (Refer to figure 6-6-20, above.)

- Erase the active startup-config file. This generates a new, default startup-config file that always results when the switch automatically reboots after deletion of the currently active startup-config file. (Refer to “Erasing a Startup-Config File” on page 6-35.)

Transitioning to Multiple Configuration Files

At the first reboot with a software release supporting multiple configuration, the switch:

- Assigns the filename **oldConfig** to the existing startup-config file (which is stored in memory slot 1).
- Saves a copy of the existing startup-config file in memory slot 2 with the filename **workingConfig**.
- Assigns the **workingConfig** file as the active configuration and the default configuration for all subsequent reboots using either primary or secondary flash.

```
ProCurve(config)# show config files
```

Configuration files:

id	act	pri	sec	name
1				oldConfig
2	*	*	*	workingConfig
3				

Figure 6-21. Switch Memory Assignments After the First Reboot from Software Supporting Multiple Configuration

In the above state, the switch always:

- Uses the **workingConfig** file to reboot

The commands described later in this section enable you to view the current multiple configuration status, manage multiple startup-config files, configure reboot policies, and override reboot policies on a per-instance basis.

Listing and Displaying Startup-Config Files

Command	Page
show config files	Below
show config < filename >	6-31

Viewing the Startup-Config File Status with Multiple Configuration Enabled

Rebooting the switch automatically enables the multiple configuration feature.

Syntax: show config files

This command displays the available startup-config files on the switch and the current use of each file.

id: *Identifies the memory slot for each startup-config file available on the switch.*

act: *An asterisk (*) in this column indicates that the corresponding startup-config file is currently in use.*

pri: *An asterisk (*) in this column indicates that the corresponding startup-config file is currently assigned to the primary boot path.*

sec: *An asterisk (*) in this column indicates that the corresponding startup-config file is currently assigned to the secondary boot path.*

name: *Shows the filename for each listed startup-config file in the switch. Refer to “Renaming an Existing Startup-Config File” on page 6-34 for the command you can use to change existing startup-config filenames.*

*In the default configuration, if the switch was shipped from the factory with software installed in both the primary and secondary boot paths, then one startup-config file named **config1** is used for both paths and is stored in memory slot 1. Memory slots 2 and 3 are empty in this default configuration.*

Displaying the Content of A Specific Startup-Config File

With Multiple Configuration enabled, the switch can have up to three startup-config files. Because the **show config** command always displays the content of the currently active startup-config file, the command extension shown below is needed to allow viewing the contents of any other startup-config files stored in the switch.

Syntax: show config < filename >

*This command displays the content of the specified startup-config file in the same way that the **show config** command displays the content of the default (currently active) startup-config file.*

Changing or Overriding the Reboot Configuration Policy

Command	Page
startup-default [primary secondary] config < filename >	Below
boot system flash < primary secondary > config < filename >	6-33

You can boot the switch using any available startup-config file.

Changing the Reboot Configuration Policy. For a given reboot, the switch automatically reboots from the startup-config file assigned to the flash location (primary or secondary) being used for the current reboot. For example, when you first download a software version that supports multiple configuration files and boot from the flash location of this version, the switch copies the existing startup-config file (named **oldConfig**) into memory slot 2, renames this file to **workingConfig**, and assigns **workingConfig** as:

- The active configuration file
- The configuration file to use when booting from either primary or secondary flash.

In this case, the switch is configured to automatically use the **workingConfig** file in memory slot 2 for all reboots.

You can use the following command to change the current policy so that the switch automatically boots using a different startup-config file.

Syntax: startup-default [primary | secondary] config < filename >

Specifies a boot configuration policy option:

[primary | secondary] config < filename >: *Designates the startup-config file to use in a reboot with the software version stored in a specific flash location. Use this option to change the reboot policy for either primary or secondary flash, or both.*

config < filename >: *Designates the startup-config file to use for all reboots, regardless of the flash version used. Use this option when you want to automatically use the same startup-config file for all reboots, regardless of the flash source used.*

For redundant management systems, this command affects both the active management module and the standby management module. The config file is copied immediately to the standby management module and becomes the default on that module when the next bootup occurs, unless redundancy is disabled or the standby module has failed selftest.

Note: *To override the current reboot configuration policy for a single reboot instance, use the **boot system flash** command with the options described under “Overriding the Default Reboot Configuration Policy” on page 6-33.*

For example, suppose:

- Software release “A” is stored in primary flash and a later software release is stored in secondary flash.
- The system operator is using memory slot 1 for a reliable, minimal configuration (named **minconfig**) for the software version in the primary flash, and slot 2 for a modified startup-config file (named **newconfig**) that includes untested changes for improved network operation with the software version in secondary flash.

The operator wants to ensure that in case of a need to reboot by pressing the Reset button, or if a power failure occurs, the switch will automatically reboot with the minimal startup-config file in memory slot 1. Since a reboot due to pressing the Reset button or to a power cycle always uses the software version in primary flash, the operator needs to configure the switch to always boot from primary flash with the startup-config file named **minconfig** (in memory slot 1). Also, whenever the switch boots from secondary flash, the operator also wants the startup-config named **newconfig** to be used. The following two commands configure the desired behavior.


```
ProCurve(config)# startup-default pri config minconfig  
ProCurve(config) # startup-default sec config newconfig.
```

Overriding the Default Reboot Configuration Policy. This command provides a method for manually rebooting with a specific startup-config file other than the file specified in the default reboot configuration policy.

Syntax: boot system flash < primary | secondary > config < filename >

Specifies the name of the startup-config file to apply for the immediate boot instance only. This command overrides the current reboot policy.

Using Reload To Reboot From the Current Flash Image and Startup-Config File.

Syntax: reload

*This command boots the switch from the currently active flash image and startup-config file. Because **reload** bypasses some subsystem self-tests, the switch boots faster than if you use a **boot** command.*

Note: To identify the currently active startup-config file, use the **show config files** command.

Managing Startup-Config Files in the Switch

Command	Page
rename config < current-filename > < newname-str >	6-34
copy config < source-filename > config < dest-filename >	6-34
erase config < filename > startup-config	6-35
Erase startup-config using the front-panel Clear + Reset Buttons	6-37

Renaming an Existing Startup-Config File

Syntax: rename config < current-filename > < newname-str >

This command changes the name of an existing startup-config file. A file name can include up to 63, alphanumeric characters. Blanks are allowed in a file name enclosed in quotes (“ ” or ‘ ’). (File names are not case-sensitive.)

For redundant management systems, renaming a config file affects both the active management module and the standby management module, unless redundancy is disabled or the standby module failed selftest.

Creating a New Startup-Config File

The switch allows up to three startup-config files. You can create a new startup-config file if there is an empty memory slot or if you want to replace one startup-config file with another.

Syntax: copy config < source-filename > config < target-filename >

This command makes a local copy of an existing startup-config file by copying the contents of an existing startup-config file in one memory slot to a new startup-config file in another, empty memory slot. This enables you to use a separate configuration file to experiment with configuration changes, while preserving the source file unchanged. It also simplifies a transition from one software version to another by enabling you to preserve the startup-config file for the earlier software version while creating a separate startup-config file for the later software version. With two such versions in place, you can easily reboot the switch with the correct startup-config file for either software version.

- *If the destination startup-config file already exists, it is overwritten by the content of the source startup-config file.*
- *If the destination startup-config file does not already exist, it will be created in the first empty configuration memory slot on the switch.*
- *If the destination startup-config file does not already exist, but there are no empty configuration memory slots on the switch, then a new startup-config file is not created and instead, the CLI displays the following error message:*

Unable to copy configuration to “< target-filename >” .

For example, suppose both primary and secondary flash memory contain software release “A” and use a startup-config file named **config1**:

```
ProCurve(config)# show config files

Configuration files:

id | act pri sec | name
-----
 1 | * * * | config1
 2 |
 3 |
```

Figure 6-22. Example of Using One Startup-Config File for Both Primary and Secondary Flash

If you wanted to experiment with configuration changes to the software version in secondary flash, you could create and assign a separate startup-config file for this purpose.

```
ProCurve(config)# copy config config1 config config2
ProCurve(config)# startup-default secondary config config2
ProCurve(config)# show config files

Configuration files:

id | act pri sec | name
-----
 1 | * * * | config1
 2 | * * * | config2
 3 |
```

The first two commands copy the **config1** startup-config file to **config2**, and then make **config2** the default startup-config file for booting from secondary flash.

Figure 6-23. Example of Creating and Assigning a New Startup-Config File

Note

You can also generate a new startup-config file by booting the switch from a flash memory location from which you have erased the currently assigned startup-config file. Refer to “Erasing a Startup-Config File” in the next section.

Erasing a Startup-Config File

You can erase any of the startup-config files in the switch’s memory slots. In some cases, erasing a file causes the switch to generate a new, default-configuration file for the affected memory slot.

In a redundant management system, this command erases the config or startup config file on both the active and the standby management modules as long as redundancy has not been disabled. If the standby management module is not in standby mode or has failed selftest, the config or startup config file is not erased.

Syntax: erase < config < filename >> | startup-config >

config < filename >: *This option erases the specified startup-config file. If the specified file is not the currently active startup-config file, then the file is simply deleted from the memory slot it occupies. If the specified file is the currently active startup-config file, then the switch creates a new, default startup-config file with the same name as the erased file, and boots using this file. (This new startup-config file contains only the default configuration for the software version used in the reboot.)*

Note: *Where a file is assigned to either the primary or the secondary flash, but is not the currently active startup-config file, erasing the file does not remove the flash assignment from the memory slot for that file. Thus, if the switch boots using a flash location that does not have an assigned startup-config, then the switch creates a new, default startup-config file and uses this file in the reboot. (This new startup-config file contains only the default configuration for the software version used in the reboot.) Executing **write memory** after the reboot causes a switch-generated filename of **configx** to appear in the **show config files** display for the new file, where **x** corresponds to the memory slot number.*

startup-config: *This option erases the currently active startup-config file and reboots the switch from the currently active flash memory location. The erased startup-config file is replaced with a new startup-config file. The new file has the same filename as the erased file, but contains only the default configuration for the software version in the flash location (primary or secondary) used for the reboot. For example, suppose the last reboot was from primary flash using a configuration file named **minconfig**. Executing **erase startup-config** replaces the current content of **minconfig** with a default configuration and reboots the switch from primary flash.*

Figure 6-24 illustrates using **erase config < filename >** to remove a startup-config file.

```
ProCurve(config)# show config files
Configuration files:
id | act pri sec | name
-----
 1 | * * * | minconfig
 2 | * * * | config2
 3 | * * * | config3

ProCurve(config)# erase config config3
ProCurve(config)# show config files
Configuration files:
id | act pri sec | name
-----
 1 | * * * | minconfig
 2 | * * * | config2
 3 | * * * |
```

Figure 6-24. Example of Erasing a Non-Active Startup-Config File

With the same memory configuration as is shown in the bottom portion of figure 6-24, executing **erase startup-config** boots the switch from primary flash, resulting in a new file named **minconfig** in the same memory slot. The new file contains the default configuration for the software version currently in primary flash.

Using the Clear + Reset Button Combination To Reset the Switch to Its Default Configuration

The Clear + Reset button combination described in the *Installation and Getting Started Guide* produces these results. That is, when you press the Clear + Reset button combination, the switch:

- Overwrites the content of the startup-config file currently in memory slot 1 with the default configuration for the software version in primary flash, and renames this file to **config1**.
- Erases any other startup-config files currently in memory.
- Configures the new file in memory slot 1 as the default for both primary and secondary flash locations (regardless of the software version currently in secondary flash).
- Boots the switch from primary flash using the new startup-config file.

```
ProCurve Switch 5304XL# sho config files
```

Configuration files:				
id	act	pri	sec	name
1	*	*	*	config1
2				
3				

Pressing Clear + Reset:

- Replaces all startup-config files with a single file named **config1** that contains the default configuration for the software version in primary flash.
- Resets the Active, Primary, and Secondary assignments as shown here.

Figure 6-25. Example of Clear + Reset Result

Transferring Startup-Config Files To or From a Remote Server

Command	Page
copy config < src-file > tftp < ip-addr > < remote-file > < pc unix > [oobm]	below
copy tftp config < dest-file > < ip-addr > < remote-file > < pc unix > [oobm]	below
copy config < src-file > xmodem < pc unix >	6-40
copy xmodem config < dest-file > < pc unix >	6-40

TFTP: Copying a Configuration File to a Remote Host

Syntax: copy config < src-file > tftp < ip-addr > < remote-file > < pc | unix > [oobm]

This is an addition to the copy tftp command options. Use this command to upload a configuration file from the switch to a TFTP server.

*For switches that have a separate out-of-band management port, the **oobm** parameter specifies that the TFTP traffic will go out through the out-of-band management interface. If this parameter is not specified, the TFTP traffic goes out through the data interface. The **oobm** parameter is not available on switches that do not have a separate out-of-band management port. Refer to Appendix I, “Network Out-of-Band Management” in this guide for more information on out-of-band management.*

For more on using TFTP to copy a file to a remote server, refer to “TFTP: Copying a Configuration File to a Remote Host” on page A-30.

For example, the following command copies a startup-config file named **test-01** from the switch to a (UNIX) TFTP server at IP address 10.10.28.14:

```
ProCurve(config)# copy config test-01 tftp 10.10.28.14  
test-01.txt unix
```

TFTP: Copying a Configuration File from a Remote Host

Syntax: copy tftp config < dest-file > < ip-addr > < remote-file > < pc | unix > [oobm]

This is an addition to the copy tftp command options. Use this command to download a configuration file from a TFTP server to the switch.

*For switches that have a separate out-of-band management port, the **oobm** parameter specifies that the TFTP traffic must come in through the out-of-band management interface. If this parameter is not specified, the TFTP traffic comes in through the data interface. The **oobm** parameter is not available on switches that do not have a separate out-of-band management port. Refer to Appendix I, “Network Out-of-Band Management” in this guide for more information on out-of-band management.*

Note: This command requires an empty memory slot in the switch. If there are no empty memory slots, the CLI displays the following message:

Unable to copy configuration to "< filename >".

*You can erase one or more configuration files using the **erase config <filename>** command.*

For more on using TFTP to copy a file from a remote host, refer to “TFTP: Copying a Configuration File from a Remote Host” on page A-31.

For example, the following command copies a startup-config file named **test-01.txt** from a (UNIX) TFTP server at IP address 10.10.28.14 to the first empty memory slot in the switch:

```
ProCurve(config)# copy tftp config test-01 10.10.28.14  
test-01.txt unix
```

Xmodem: Copying a Configuration File to a Serially Connected Host

Syntax: copy config < filename > xmodem < pc | unix >

*This is an addition to the **copy < config > xmodem** command options. Use this command to upload a configuration file from the switch to an Xmodem host.*

For more on using Xmodem to copy a file to a serially connected host, refer to “Xmodem: Copying a Configuration File to a Serially Connected PC or UNIX Workstation” on page A-33.

Xmodem: Copying a Configuration from a Serially Connected Host

Syntax: copy xmodem config < dest-file > < pc | unix >

*This is an addition to the **copy xmodem** command options. Use this command to download a configuration file from an Xmodem host to the switch.*

For more on using Xmodem to copy a file from a serially connected host, refer to “Xmodem: Copying a Configuration File from a Serially Connected PC or UNIX Workstation” on page A-33.

Operating Notes for Multiple Configuration Files

- SFTP/SCP: The configuration files are available for sftp/scp transfer as **/ctg/< filename >**.

Automatic Configuration Update with DHCP Option 66

ProCurve switches are initially booted up with the factory-shipped configuration file. This feature provides a way to automatically download a different configuration file from a TFTP server using DHCP Option 66. The prerequisites for this to function correctly are:

- One or more DHCP servers with Option 66 are enabled
- One or more TFTP servers has the desired configuration file.

Caution

This feature must use configuration files generated on the switch to function correctly. If you use configuration files that were not generated on the switch, and then enable this feature, the switch may reboot continuously.

CLI Command

The command to enable the configuration update using Option 66 is:

Syntax: [no] dhcp config-file-update

Enables configuration file update using Option 66.

Default: Enabled

```
ProCurve(config)# dhcp config-file-update
```

Figure 6-26. Example of Enabling Configuration File Update Using Option 66

Possible Scenarios for Updating the Configuration File

The following table shows various network configurations and how Option 66 is handled.

Scenario	Behavior
Single Server serving Multiple VLANs	<ul style="list-style-type: none">• Each DHCP-enabled VLAN interface initiates DHCPDISCOVER message, receives DHCPOFFER from the server, and send DHCPREQUEST to obtain the offered parameters.• If multiple interfaces send DHCPREQUESTs, it's possible that more than one DHCPACK is returned with a valid Option 66.• Evaluating and updating the configuration file occurs only on the primary VLAN.• Option 66 is ignored by any interfaces not belonging to the primary VLAN.
Multiple Servers serving a Single VLAN	<ul style="list-style-type: none">• Each DHCP-enabled VLAN interface initiates one DHCPDISCOVER and receives one or more DHCPOFFER messages.• Each interface accepts the best offer.• Option 66 is processed only for the interface belonging to the primary VLAN.
Multiple Servers serving Multiple VLANs	<ul style="list-style-type: none">• Each DHSP-enabled VLAN interface initiates DHCPDISCOVER and receives one or more DHCPOFFER messages.• Each interface accepts the best offer.• Option 66 is processed only for the interface belonging to the primary VLAN.
Multi-homed Server serving Multiple VLANs	<ul style="list-style-type: none">• The switch perceives the multi-homed server as multiple separate servers.• Each DHCP-enabled VLAN interface initiates DHCPDISCOVER and receives one DHCPOFFER message.• Each interface accepts the offer.• Option 66 is processed only for the interface belonging to the primary VLAN.

Operating Notes

Replacing the Existing Configuration File: After the DHCP client downloads the configuration file, the switch compares the contents of that file with the existing configuration file. If the content is different, the new configuration file replaces the existing file and the switch reboots.

Option 67 and the Configuration File Name: Option 67 includes the name of the configuration file. If the DHCPACK contains this option, it overrides the default name for the configuration file (switch.cfg)

Global DHCP Parameters: Global parameters are processed only if received on the primary VLAN.

Best Offer: The “Best Offer” is the best DHCP or BootP offer sent by the DHCP server in response to the DHCPREQUEST sent by the switch. The criteria for selecting the “Best Offer” are:

- DHCP is preferred over BootP
- If two BootP offers are received, the first one is selected
- For two DHCP offers:
 - The offer from an authoritative server is selected
 - If there is no authoritative server, the offer with the longest lease is selected

Log Messages

The file transfer is implemented by the existing TFTP module. The system logs the following message if an incorrect IP address is received for Option 66:

“Invalid IP address <ip-address> received for DHCP Option 66”

Interface Access and System Information

Contents

Overview	7-2
Interface Access: Console/Serial Link, Web, and Inbound Telnet .	7-3
Menu: Modifying the Interface Access	7-4
CLI: Modifying the Interface Access	7-5
Denying Interface Access by Terminating Remote Management Sessions	7-11
System Information	7-12
Menu: Viewing and Configuring System Information	7-13
CLI: Viewing and Configuring System Information	7-14
Web: Configuring System Parameters	7-18

Overview

This chapter describes how to:

- View and modify the configuration for switch interface access
- Use the CLI **kill** command to terminate a remote session
- View and modify switch system information

For help on how to actually use the interfaces built into the switch, refer to:

- Chapter 3, “Using the Menu Interface”
- Chapter 4, “Using the Command Line Interface (CLI)”
- Chapter 5, “Using the ProCurve Web Browser Interface”

Why Configure Interface Access and System Information? The interface access features in the switch operate properly by default. However, you can modify or disable access features to suit your particular needs. Similarly, you can choose to leave the system information parameters at their default settings. However, modifying these parameters can help you to more easily distinguish one device from another in your network.

Interface Access: Console/Serial Link, Web, and Inbound Telnet

Interface Access Features

Feature	Default	Menu	CLI	Web
Inactivity Time	0 Minutes (disabled)	page 7-4	page 7-9	—
Inbound Telnet Access	Enabled	page 7-4	page 7-5	—
Outbound Telnet Access	n/a	—	page 7-6	—
Web Browser Interface Access	Enabled	page 7-4	page 7-8	—
Terminal type	VT-100	—	page 7-9	—
Event Log event types to list (Displayed Events)	All	—	page 7-9	—
Baud Rate	Speed Sense	—	page 7-9	—
Flow Control	XON/XOFF	—	page 7-9	—

In most cases, the default configuration is acceptable for standard operation.

Note

Basic switch security is through passwords. You can gain additional security by using the security features described in the Access Security Guide for your switch. You can also simply block unauthorized access via the web browser interface or Telnet (as described in this section) and installing the switch in a locked environment.

Menu: Modifying the Interface Access

The menu interface enables you to modify these parameters:

- Inactivity Timeout
- Inbound Telnet Enabled
- Web Agent Enabled

To Access the Interface Access Parameters:

1. From the Main Menu, Select...

2. Switch Configuration...

1. System Information

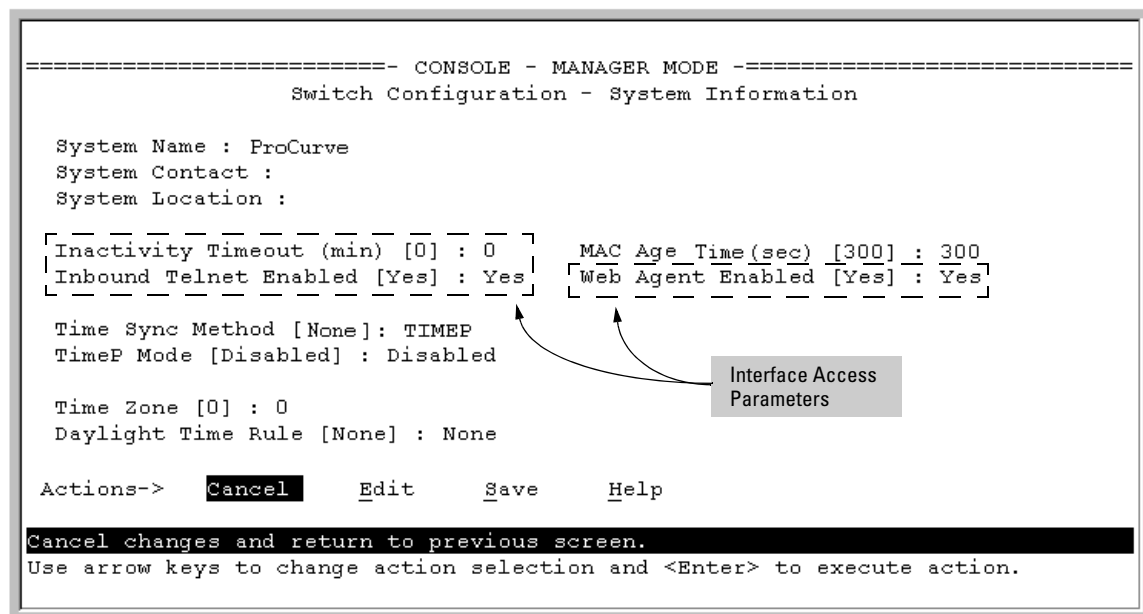


Figure 7-1. The Default Interface Access Parameters Available in the Menu Interface

2. Press [E] (for Edit). The cursor moves to the **System Name** field.
3. Use the arrow keys (J, U, L, R) to move to the parameters you want to change.

Refer to the online help provided with this screen for further information on configuration options for these features.

4. When you have finished making changes to the above parameters, press [Enter], then press [S] (for Save).

CLI: Modifying the Interface Access

Interface Access Commands Used in This Section

show console	below
[no] telnet-server	below
[no] web-management	page 7-8
console	page 7-9

Listing the Current Console/Serial Link Configuration. This command lists the current interface access parameter settings.

Syntax: show console

This example shows the switch's default console/serial configuration.

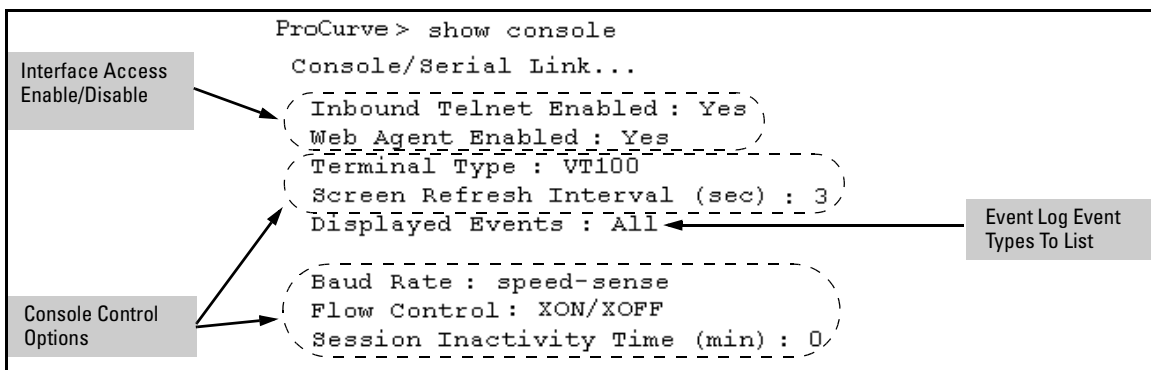


Figure 7-2. Listing of Show Console Command

Reconfigure Inbound Telnet Access. In the default configuration, inbound Telnet access is enabled.

Interface Access and System Information

Interface Access: Console/Serial Link, Web, and Inbound Telnet

Syntax: [no] telnet-server [listen <oobm | data | both>]

Enables or disables inbound Telnet access on a switch.

*Use the **no** version of the command to disable inbound Telnet access.*

*The **listen** parameter is available only on switches that have a separate out-of-band management port. Values for this parameter are:*

- **oobm** — *inbound Telnet access is enabled only on the out-of-band management port.*
- **data** — *inbound Telnet access is enabled only on the data ports.*
- **both** — *inbound Telnet access is enabled on both the out-of-band management port and on the data ports. This is the default value.*

Refer to Appendix I, “Network Out-of-Band Management” in this guide for more information on out-of-band management.

*The **listen** parameter is not available on switches that do not have a separate out-of-band management port.*

To disable inbound Telnet access:

```
ProCurve(config)# no telnet-server
```

To re-enable inbound Telnet access:

```
ProCurve(config)# telnet-server
```

Outbound Telnet to Another Device. This feature operates independently of the telnet-server status and enables you to Telnet to another device that has an IP address.

Syntax: telnet <ipv4-addr | ipv6-addr | hostname | switch-num> [oobm]

Initiates an outbound telnet session to another network device. The destination can be specified as:

- *IPv4 address*
- *IPv6 address*
- *Hostname*
- *Stack number of a member switch (1-16) if the switch is a commander in a stack and stacking is enabled*

*For switches that have a separate out-of-band management port, the **oobm** parameter specifies that the Telnet traffic will go out from the out-of-band management interface. If this parameter is not specified, the Telnet traffic goes out from the data interface. The oobm parameter is not available on switches that do not have a separate out-of-band management port. Refer to Appendix I, “Network Out-of-Band Management” in this guide for more information on out-of-band management.*

For example, if the host “Labswitch” is in the domain abc.com, you can enter the following command and the destination is resolved to “Labswitch.abc.com”.

```
ProCurve(config)# telnet Labswitch
```

You can also enter the full domain name in the command:

```
ProCurve(config)# telnet Labswitch.abc.com
```

You can use the **show telnet** command to display the resolved IP address.

Interface Access and System Information

Interface Access: Console/Serial Link, Web, and Inbound Telnet

```
ProCurve(config)# show telnet

Telnet Activity

-----
Session : ** 1
Privilege: Manager
From    : Console
To      :

-----
Session : ** 2
Privilege: Manager
From    : 12.13.14.10
To      : 15.33.66.20

-----
Session : ** 3
Privilege: Operator
From    : 2001:db7:5:0:203:4ff:fe0a:251
To      : 2001:db7:5:0:203:4ff1:fddd:12
```

Figure 7-3. Example of show telnet Command Displaying Resolved IP Addresses

Reconfigure Web Browser Access. In the default configuration, web browser access is enabled.

Syntax: [no] web-management [[listen <oobm | data | both>]

*Use the **no** version of the command to disable inbound HTTP access.*

*The **listen** parameter is available only on switches that have a separate out-of-band management port. Values for this parameter are:*

- **oobm** — *inbound HTTP access is enabled only on the out-of-band management port.*
- **data** — *inbound HTTP access is enabled only on the data ports.*
- **both** — *inbound HTTP access is enabled on both the out-of-band management port and on the data ports. This is the default value.*

Refer to Appendix I, “Network Out-of-Band Management” in this guide for more information on out-of-band management.

*The **listen** parameter is not available on switches that do not have a separate out-of-band management port.*

To disable web browser access:

```
ProCurve(config)# no web-management
```

To re-enable web browser access:

```
ProCurve(config)# web-management
```

Reconfigure the Console/Serial Link Settings. You can reconfigure one or more console parameters with one console command.

Syntax: console

```
[terminal < vt100 | ansi | none >]
[screen-refresh < 1 | 3 | 5 | 10 | 20 | 30 | 45 | 60 >]
[baud-rate
  < speed-sense | 1200 | 2400 | 4800 | 9600 | 19200 | 38400 | 57600 |
  1155200 >]
[ flow-control < xon/xoff | none >]
[inactivity-timer < 0 | 1 | 5 | 10 | 15 | 20 | 30 | 60 | 120 >]
[events < none | all | non-info | critical | debug]
[local-terminal < vt 100 | none | ansi >]
```

Note

If you change the Baud Rate or Flow Control settings for the switch, you should make the corresponding changes in your console access device. Otherwise, you may lose connectivity between the switch and your terminal emulator due to differences between the terminal and switch settings for these two parameters.

All console parameter changes except **events** and **inactivity-timer** require that you save the configuration with **write memory** and then execute **boot** before the new console configuration will take effect.

For example, to use one command to configure the switch with the following:

- VT100 operation
- 19,200 baud
- No flow control
- 10-minute inactivity time
- Critical log events

you would use the following command sequence:

Interface Access and System Information

Interface Access: Console/Serial Link, Web, and Inbound Telnet

```
ProCurve(config)# console terminal vt100 baud-rate 19200 flow-control none
inactivity-timer 10 events critical
Command will take effect after saving configuration and reboot.
ProCurve(config)# write memory
ProCurve(config)# reload
```

The switch implements the Event Log change immediately. The switch implements the other console changes after executing **write memory** and **reload**.

Figure 7-4. Example of Executing the Console Command with Multiple Parameters

Note

When using redundant management, console settings, such as mode, flow-control and baud-rate, are the same on both management modules. There cannot be individual settings for each management module.

You can also execute a series of console commands and then save the configuration and boot the switch. For example:

```
Configure the individual parameters.
ProCurve(config)# console baud-rate speed-sense
Command will take effect after saving configuration and reboot
ProCurve(config)# console flow-control xon/xoff
Command will take effect after saving configuration and reboot
Save the changes.
ProCurve(config)# console inactivity-timer 0
Command will take effect after saving configuration and reboot
Boot the switch.
ProCurve(config)# write memory
ProCurve(config)# reload
```

Figure 7-5. Example of Executing a Series of Console Commands

Denying Interface Access by Terminating Remote Management Sessions

The switch supports up to five management sessions. You can use **show ip ssh** to list the current management sessions, and **kill** to terminate a currently running remote session. (**Kill** does not terminate a Console session on the serial port, either through a direct connection or via a modem. It does not affect the console on the standby module.)

Syntax: kill [< session-number >]

For example, if you are using the switch's serial port for a console session and want to terminate a currently active Telnet session, you would do the following:

```
ProCurve(config)# show ip ssh
SSH Enabled           : Yes

IP Port Number        : 22
Timeout (sec)         : 120
Server Key Size (bits) : 512

Ses Type      Source IP and Port
-----
1  console
2  telnet
3  ssh      15.30.252.195:1531
4  inactive
5  inactive

ProCurve(config)# kill 2
ProCurve(config)# show ip ssh
SSH Enabled           : Yes

IP Port Number        : 22
Timeout (sec)         : 120
Server Key Size (bits) : 512

Ses Type      Source IP and Port
-----
1  console
2  inactive
3  ssh      15.30.252.195:1531
4  inactive
5  inactive
```

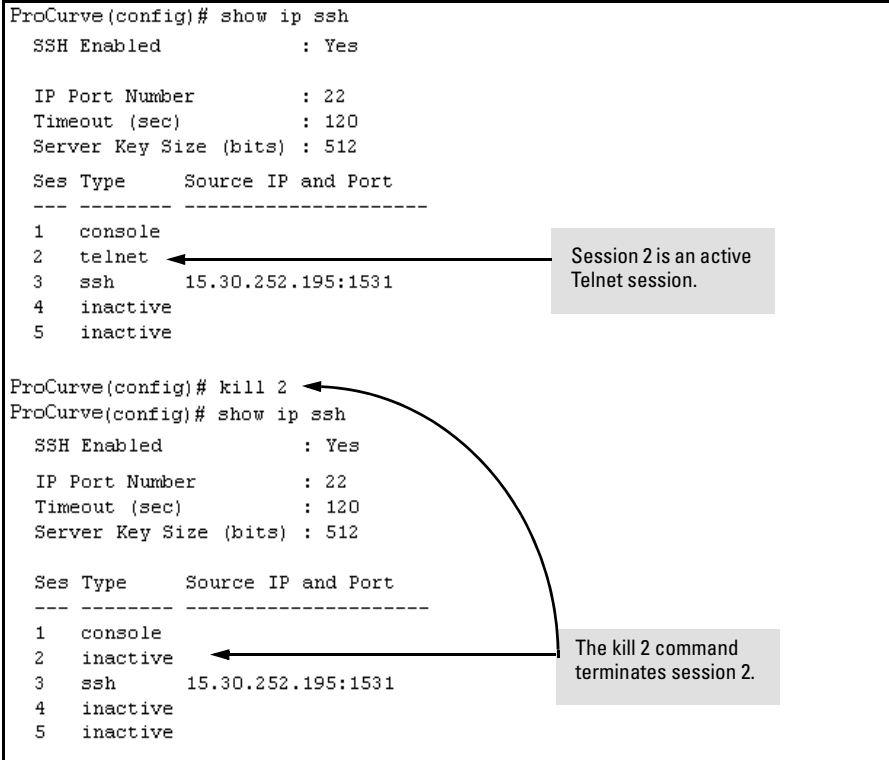


Figure 7-6. Example of Using the “Kill” Command To Terminate a Remote Session

System Information

System Information Features

Feature	Default	Menu	CLI	Web
System Name	<i>switch product name</i>	page 7-13	page 7-15	page 7-18
System Contact	n/a	page 7-13	page 7-15	page 7-18
System Location	n/a	page 7-13	page 7-15	page 7-18
MAC Age Time	300 seconds	page 7-13	page 7-17	—
Time Sync Method	None	See Chapter 9, “Time Protocols”.		
Time Zone	0	page 7-13	page 7-17	—
Daylight Time Rule	None	page 7-13	page 7-17	—
Time	January 1, 1990 at 00:00:00 at last power reset	—	page 7-18	—

Configuring system information is optional, but recommended.

System Name: Using a unique name helps you to identify individual devices where you are using an SNMP network management tool such as ProCurve Manager.

System Contact and Location: This information is helpful for identifying the person administratively responsible for the switch and for identifying the locations of individual switches.

MAC Age Time: The number of seconds a MAC address the switch has learned remains in the switch’s address table before being aged out (deleted). Aging out occurs when there has been no traffic from the device belonging to that MAC address for the configured interval.

Time Sync Method: Selects the method (TimeP or SNTP) the switch will use for time synchronization. For more on this topic, refer to Chapter 9, “Time Protocols”.

Time Zone: The number of minutes your time zone location is to the West (+) or East (-) of Coordinated Universal Time (formerly GMT). The default **0** means no time zone is configured. For example, the time zone for Berlin, Germany is + 60 (minutes) and the time zone for Vancouver, Canada is - 480 (minutes).

Daylight Time Rule: Specifies the daylight savings time rule to apply for your location. The default is **None**. (For more on this topic, refer to Appendix D, “Daylight Savings Time on ProCurve Switches.”)

Time: Used in the CLI to specify the time of day, the date, and other system parameters.

Menu: Viewing and Configuring System Information

To access the system information parameters:

1. From the Main Menu, Select...
 2. **Switch Configuration...**
 1. **System Information**

```
----- CONSOLE - MANAGER MODE -----
Switch Configuration - System Information

System Name : ProCurve
System Contact :
System Location :

Inactivity Timeout (min) [0] : 0
Inbound Telnet Enabled [Yes] : Yes

Time Sync Method [None] : TIMEP
TimeP Mode [Disabled] : Disabled

Time Zone [0] : 0
Daylight Time Rule [None] : None

Actions->  Cancel  Edit  Save  Help

Cancel changes and return to previous screen.
Use arrow keys to change action selection and <Enter> to execute action.
```

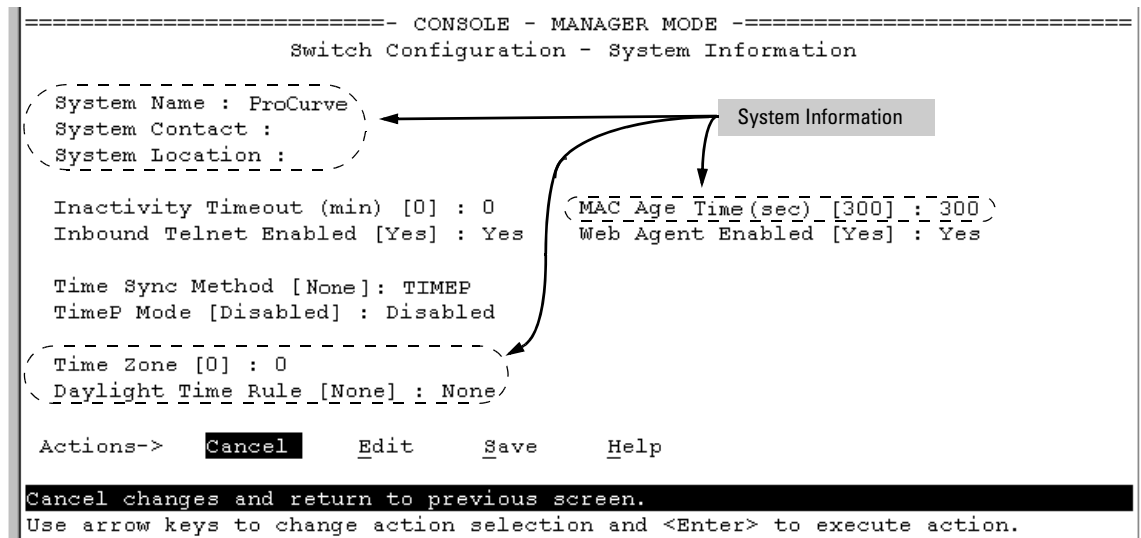


Figure 7-7. The System Information Configuration Screen (Default Values)

Note

To help simplify administration, it is recommended that you configure **System Name** to a character string that is meaningful within your system.

2. Press **[E]** (for **E**dit). The cursor moves to the **System Name** field.
3. Refer to the online help provided with this screen for further information on configuration options for these features.
4. When you have finished making changes to the above parameters, press **[Enter]**, then press **[S]** (for **S**ave) and return to the Main Menu.

CLI: Viewing and Configuring System Information

System Information Commands Used in This Section

show system information	below
hostname	below
snmp-server [contact] [location]	below
mac-age-time	page 7-17
time	
timezone	page 7-17
daylight-time-rule	page 7-17
date	page 7-18
time	

Listing the Current System Information. This command lists the current system information settings.

Syntax: show system information

This example shows the switch's default console configuration.

```
ProCurve# show system information

Status and Counters - General System Information

System Name       : ProCurve
System Contact    :
System Location   :

MAC Age Time (sec) : 300

Time Zone         : 0
Daylight Time Rule : None
```

Figure 7-8. Example of CLI System Information Listing

Configure a System Name, Contact, and Location for the Switch. To help distinguish one switch from another, configure a plain-language identity for the switch.

Syntax: `hostname < name-string >`
`snmp-server [contact <system-contact>] [location <system-location>]`

Each field allows up to 255 characters.

For example, to name the switch “Blue” with “Next-4474” as the system contact, and “North-Data-Room” as the location:

```
ProCurve(config)# hostname Blue
Blue(config)# snmp-server contct Ext-4474 location North-Data-Room
Blue(config)# show system-information

Status and Counters - General System Information
-----
System Name       : Blue
System Contact    : Ext-4474
System Location   : North-Data-Room

MAC Age Time (sec) : 300

Time Zone         : 0
Daylight Time Rule : None

Firmware revision : E.08.30      Base MAC Addr   : 0001e7-a0ec00
ROM Version        : E.05.04      Serial Number    : S000394041

Up Time           : 14 mins      Memory - Total  : 25,038,312
CPU Util (%)      : 1             Free           : 20,087,448

IP Mgmt - Pkts Rx : 0             Packet - Total  : 832
          Pkts Tx : 0             Buffers  Free   : 783
                                   Lowest      : 768

-- MORE --, next page: Space, next line: Enter, quit: Control-C
```

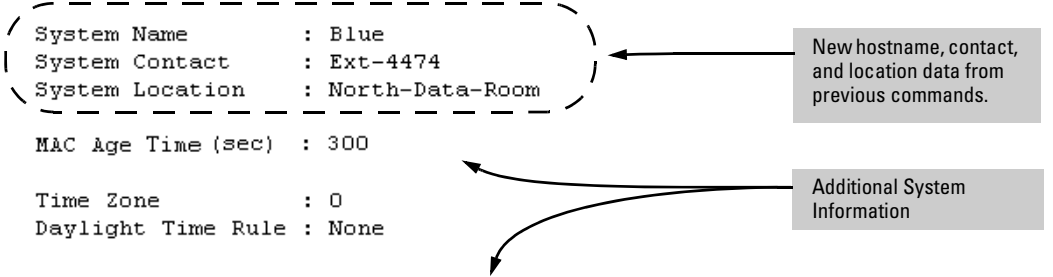


Figure 7-9. System Information Listing After Executing the Preceding Commands

The menu interface will only display up to 47 characters although you can specify a name up to 255 characters in length. A message beginning with “+” displays if the name exceeds 47 characters. You can use the CLI **show running**, **show config**, or **show system information** commands to see the complete text. The menu interface is shown in Figure 7-10.

Interface Access and System Information
System Information

```
MENU
ProCurve Switch 5406zl                               24-Oct-2006  12:41:47
===== TELNET - MANAGER MODE =====
                Switch Configuration - System Information

System Name : Blue Switch
System Contact : Bill_Smith
System Location : + characters of the location are missing. It's too long.

Inactivity Timeout (min) [0] : 0           MAC Age Time (sec) [300] : 300
Inbound Telnet Enabled [Yes] : Yes         Web Agent Enabled [Yes] : Yes
Time Sync Method [None] : TIMEP
TimeP Mode [Disabled] : Disabled

Tftp-enable [Yes] : Yes
Time Zone [0] : 0
Daylight Time Rule [None] : None

Actions->   Cancel      Edit      Save      Help

Cancel changes and return to previous screen.
Use arrow keys to change action selection and <Enter> to execute action.
```

Figure 7-10. Menu Screen Showing System Information

The Web Browser interface also allows you to enter a maximum of 255 characters. You can view all the characters by using the cursor to scroll through the field.

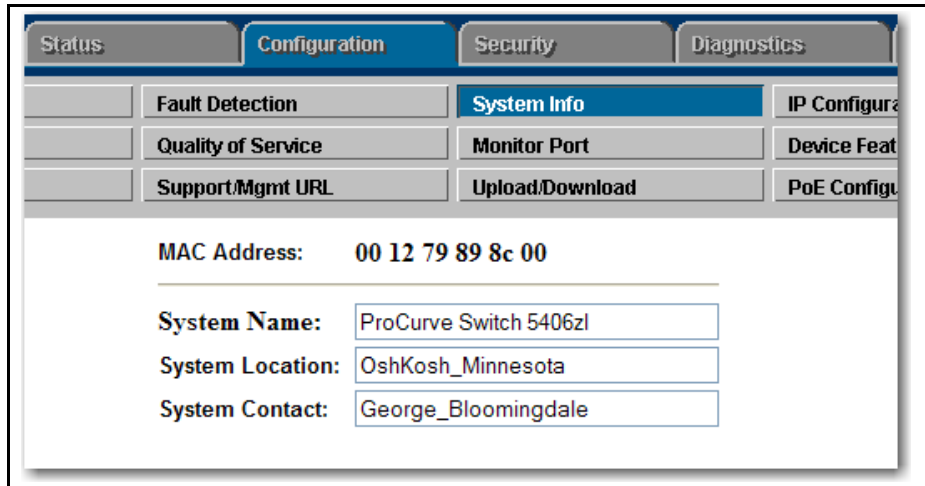


Figure 7-11. System Location and System Contact in the Web Browser

Reconfigure the MAC Age Time for Learned MAC Addresses. This command corresponds to the MAC Age Interval in the menu interface, and is expressed in seconds.

Syntax: `mac-age-time < 10 - 1000000 > (seconds)`

*Allows you to set the MAC address table's **age-out** interval. An address is aged out if the switch does not receive traffic from that MAC address for the **age-out** interval, measured in seconds.
Default: 300 seconds.*

For example, to configure the age time to seven minutes:

```
ProCurve(config)# mac-age-time 420
```

Configure the Time Zone and Daylight Time Rule. These commands:

- Set the time zone you want to use
- Define the daylight time rule for keeping the correct time when daylight-saving-time shifts occur.

Syntax: `time timezone < -720 - 840 >`

`time daylight-time-rule < none | alaska | continental-us-and-canada | middle-europe-and-portugal | southern-hemisphere | western-europe | user-defined >`

East of the 0 meridian, the sign is "+". West of the 0 meridian, the sign is "-".

For example, the time zone setting for Berlin, Germany is +60 (zone +1, or 60 minutes), and the time zone setting for Vancouver, Canada is -480 (zone -8, or -480 minutes). To configure the time zone and daylight time rule for Vancouver, Canada:

```
ProCurve(config)# time timezone -480  
daylight-time-rule continental-us-and-canada
```

Configure the Time and Date. The switch uses the time command to configure both the time of day and the date. Also, executing time without parameters lists the switch's time of day and date. Note that the CLI uses a 24-hour clock scheme; that is, hour (*hh*) values from 1 p.m. to midnight are input as 13 - 24, respectively.

Syntax: time [*hh:mm* [:*ss*]] [*mm/dd* / [*yy*] *yy*]

For example, to set the switch to 9:45 a.m. on November 17, 2002:

```
ProCurve(config)# time 9:45 11/17/02
```

Note

Executing **reload** or **boot** resets the time and date to their default startup values.

Web: Configuring System Parameters

In the web browser interface, you can enter the following system information:

- System Name
- System Location
- System Contact

For access to the MAC Age Interval and the Time parameters, use the menu interface or the CLI.

Configure System Parameters in the Web Browser Interface.

1. Click on the **Configuration** tab.
2. Click on **[System Info]**.
3. Enter the data you want in the displayed fields.
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.

Configuring IP Addressing

Contents

Overview	8-2
IP Configuration	8-2
Just Want a Quick Start with IP Addressing?	8-4
IP Addressing with Multiple VLANs	8-4
Menu: Configuring IP Address, Gateway, and Time-To-Live (TTL) ..	8-5
CLI: Configuring IP Address, Gateway, and Time-To-Live (TTL)	8-7
Web: Configuring IP Addressing	8-11
How IP Addressing Affects Switch Operation	8-11
DHCP/Bootp Operation	8-12
Network Preparations for Configuring DHCP/Bootp	8-15
Loopback Interfaces	8-16
Introduction	8-16
Configuring a Loopback Interface	8-17
Displaying Loopback Interface Configurations	8-18
IP Preserve: Retaining VLAN-1 IP Addressing Across Configuration File Downloads	8-21
Operating Rules for IP Preserve	8-21
Enabling IP Preserve	8-22
Configuring a Single Source IP Address	8-25
Overview	8-25
Specifying the Source IP Address	8-25
The Source IP Selection Policy	8-26
Displaying the Source IP Interface Information	8-29
Error Messages	8-32

Overview

You can configure IP addressing through all of the switch's interfaces. You can also:

- Easily edit a switch configuration file to allow downloading the file to multiple switches without overwriting each switch's unique gateway and VLAN 1 IP addressing.
- Assign up to 32 IP addresses to a VLAN (multinetting).
- Select an IP address to use as the source address for all outgoing traffic generated by a specified software application on the switch. This allows unique identification of the software application on the server site regardless of which local interface has been used to reach the destination server.

Why Configure IP Addressing? In its factory default configuration, the switch operates as a multiport learning bridge with network connectivity provided by the ports on the switch. However, to enable specific management access and control through your network, you will need IP addressing. Table 8-1 on page 8-12 shows the switch features that depend on IP addressing to operate.

IP Configuration

IP Configuration Features

Feature	Default	Menu	CLI	Web
IP Address and Subnet Mask	DHCP/Bootp	page 8-5	page 8-7	page 8-11
Multiple IP Addresses on a VLAN	n/a	—	page 8-9	—
Default Gateway Address	none	page 8-5	page 8-7	page 8-11
Packet Time-To-Live (TTL)	64 seconds	page 8-5	page 8-7	—
Time Server (Timep)	DHCP	page 8-5	page 8-7	—
Single Source IP Addressing	outgoing IP address	—	page 8-25	—

IP Address and Subnet Mask. Configuring the switch with an IP address expands your ability to manage the switch and use its features. By default, the switch is configured to automatically receive IP addressing on the default VLAN from a DHCP/Bootp server that has been configured correctly with information to support the switch. (Refer to “DHCP/Bootp Operation” on page 8-12 for information on setting up automatic configuration from a server.) However, if you are not using a DHCP/Bootp server to configure IP addressing, use the menu interface or the CLI to manually configure the initial IP values. After you have network access to a device, you can use the web browser interface to modify the initial IP configuration if needed.

For information on how IP addressing affects switch operation, refer to “How IP Addressing Affects Switch Operation” on page 8-11.

Multinetting: Assigning Multiple IP Addresses to a VLAN. For a given VLAN you can assign up to 32 IP addresses. This allows you to combine two or more subnets on the same VLAN, which enables devices in the combined subnets to communicate normally through the network without needing to reconfigure the IP addressing in any of the combined subnets.

Default Gateway Operation. The default gateway is required when a router is needed for tasks such as reaching off-subnet destinations or forwarding traffic across multiple VLANs. The gateway value is the IP address of the next-hop gateway node for the switch, which is used if the requested destination address is not on a local subnet/VLAN. If the switch does not have a manually-configured default gateway and DHCP/Bootp is configured on the primary VLAN, then the default gateway value provided by the DHCP or Bootp server will be used. If the switch has a manually configured default gateway, then the switch uses his gateway, even if a different gateway is received via DHCP or Bootp on the primary VLAN. This is also true for manually configured TimeP, SNTP, and Time-To-Live(TTL). (In the default configuration, VLAN 1 is the Primary VLAN.) Refer to the information on Primary VLANs in the *Advanced Traffic Management Guide* for your switch.

Packet Time-To-Live (TTL) . This parameter specifies the maximum number of routers (hops) through which a packet can pass before being discarded. Each router decreases a packet’s TTL by 1 before forwarding the packet. If decreasing the TTL causes the TTL to be 0, the router drops the packet instead of forwarding it. In most cases, the default setting (64) is adequate.

Just Want a Quick Start with IP Addressing?

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

- Enter setup at the CLI Manager level prompt.

```
ProCurve# setup
```

- Select **8. Run Setup** in the Main Menu of the menu interface.

For more on using the Switch Setup screen, refer to the *Installation and Getting Started Guide* you received with the switch.

IP Addressing with Multiple VLANs

In the factory-default configuration, the switch has one, permanent default VLAN (named DEFAULT_VLAN) that includes all ports on the switch. Thus, when only the default VLAN exists in the switch, if you assign an IP address and subnet mask to the switch, you are actually assigning the IP addressing to the DEFAULT_VLAN.

Notes

- If multiple VLANs are configured, then each VLAN can have its own IP address. This is because each VLAN operates as a separate broadcast domain and requires a unique IP address and subnet mask. A default gateway (IP) address for the switch is optional, but recommended.
- In the factory-default configuration, the default VLAN (named DEFAULT_VLAN) is the switch's *primary* VLAN. The switch uses the primary VLAN for learning the default gateway address. The switch can also learn other settings from a DHCP or Bootp server, such as (packet) Time-To-Live (TTL), and Timep or SNMP settings. (Other VLANs can also use DHCP or BootP to acquire IP addressing. However, the switch's gateway, TTL, and TimeP or SNTP values, which are applied globally, and not per-VLAN, will be acquired through the primary VLAN only, unless manually set by using the CLI, Menu, or web browser interface. (If these parameters are manually set, they will *not* be overwritten by alternate values received from a DHCP or Bootp server.) For more on VLANs, refer to the chapter titled "Static Virtual LANs" in the *Advanced Traffic Management Guide* for your switch.

- The IP addressing used in the switch should be compatible with your network. That is, the IP address must be unique and the subnet mask must be appropriate for your IP network.
 - If you change the IP address through either Telnet access or the web browser interface, the connection to the switch will be lost. You can reconnect by either restarting Telnet with the new IP address or entering the new address as the URL in your web browser.
-

Menu: Configuring IP Address, Gateway, and Time-To-Live (TTL)

Do one of the following:

- To manually enter an IP address, subnet mask, set the **IP Config** parameter to **Manual** and then manually enter the IP address and subnet mask values you want for the switch.
- To use DHCP or Bootp, use the menu interface to ensure that the **IP Config** parameter is set to **DHCP/Bootp**, then refer to “DHCP/Bootp Operation” on page 8-12.

To Configure IP Addressing.

1. From the Main Menu, Select.
 2. **Switch Configuration ...**
 5. **IP Configuration**

Notes

If multiple VLANs are configured, a screen showing all VLANs appears instead of the following screen.

The Menu interface displays the IP address for any VLAN. If you use the CLI to configure the IP address on a VLAN, use the CLI **show ip** command to list them. (Refer to “Viewing the Current IP Configuration” on page 8-7.)

For descriptions of these parameters, see the online Help for this screen.

Before using the DHCP/Bootp option, refer to “DHCP/Bootp Operation” on page 8-12.

```
=====-- CONSOLE - MANAGER MODE -----  
Switch Configuration - Internet (IP) Service  
  
Default Gateway :  
Default TTL      : 64  
  
IP Config [DHCP/Bootp] : Manual  
IP Address       : 15.30.248.184  
Subnet Mask      : 255.255.248.0  
  
Actions->  Cancel  Edit  Save  Help  
Cancel changes and return to previous screen.  
Use arrow keys to change action selection and <Enter> to execute action.
```

Figure 8-1. Example of the IP Service Configuration Screen without Multiple VLANs Configured

2. Press **[E]** (for **E**dit).
3. If the switch needs to access a router, for example, to reach off-subnet destinations, select the **Default Gateway** field and enter the IP address of the gateway router.
4. If you need to change the packet Time-To-Live (TTL) setting, select **Default TTL** and type in a value between 2 and 255.
5. To configure IP addressing, select **IP Config** and do one of the following:
 - If you want to have the switch retrieve its IP configuration from a DHCP or Bootp server, at the **IP Config** field, keep the value as **DHCP/Bootp** and go to step 8.
 - If you want to manually configure the IP information, use the Space bar to select **Manual** and use the **[Tab]** key to move to the other IP configuration fields.
6. Select the **IP Address** field and enter the IP address for the switch.
7. Select the **Subnet Mask** field and enter the subnet mask for the IP address.
8. Press **[Enter]**, then **[S]** (for **S**ave).

CLI: Configuring IP Address, Gateway, and Time-To-Live (TTL)

IP Commands Used in This Section	Page
show ip	8-7
ip address < mask-length >	8-8, 8-9
ip address /< mask-bits >	8-8, 8-9
ip default-gateway	8-10
ip ttl	8-11

Viewing the Current IP Configuration.

Syntax: show ip

This command displays the IP addressing for each VLAN configured in the switch. If only the DEFAULT_VLAN exists, then its IP configuration applies to all ports in the switch. Where multiple VLANs are configured, the IP addressing is listed per VLAN. The display includes switch-wide packet time-to-live, and (if configured) the switch's default gateway and Timep configuration.

(You can also use the **show management** command to display the IP addressing and time server IP addressing configured on the switch. Refer to figure 9-6 on page 9-11.)

For example, in the factory-default configuration (no IP addressing assigned), the switch's IP addressing appears as:

```

The Default IP Configuration
ProCurve> show ip
Internet (IP) Service
  Default Gateway :
  Default TTL    : 64
  Arp Age       : 20

  TimeP Config : DHCP      TimeP Poll Interval (min) : 720

  VLAN          | IP Config  IP Address      Subnet Mask
  -----+-----
  DEFAULT_VLAN | DHCP/Bootp
  
```

Figure 8-2. Example of the Switch's Default IP Addressing

With multiple VLANs and some other features configured, **show ip** provides additional information:

```
A Switch with IP Addressing and VLANs Configured
ProCurve> show ip
Internet (IP) Service
IP Routing : Disabled
Default Gateway : 10.28.227.1
Default TTL    : 64
VLAN          | IP Config  IP Address      Subnet Mask
-----+-----
DEFAULT_VLAN | Manual     10.28.227.101   255.255.248.0
VLAN_2       | Disabled
```

Figure 8-3. Example of Show IP Listing with Non-Default IP Addressing Configured

Configure an IP Address and Subnet Mask. The following command includes both the IP address and the subnet mask. You must either include the ID of the VLAN for which you are configuring IP addressing or go to the context configuration level for that VLAN. (If you are not using VLANs on the switch—that is, if the only VLAN is the default VLAN—then the VLAN ID is always “1”.)

Note

The default IP address setting for the DEFAULT_VLAN is **DHCP/Bootp**. On additional VLANs you create, the default IP address setting is **Disabled**.

Syntax: [no] vlan < vlan-id > ip address < ip-address / mask-length >
or
[no] vlan < vlan-id > ip address < ip-address > < mask-bits >
or
vlan < vlan-id > ip address dhcp-bootp

This example configures IP addressing on the default VLAN with the subnet mask specified in mask bits.

```
ProCurve(config)# vlan 1 ip address 10.28.227.103 255.255.255.0
```

This example configures the same IP addressing as the preceding example, but specifies the subnet mask by mask length.

```
ProCurve(config)# vlan 1 ip address 10.28.227.103/24
```

This example deletes an IP address configured in VLAN 1.

```
ProCurve (config) no vlan 1 ip address 10.28.227.103/24
```

Configure Multiple IP Addresses on a VLAN (Multinetting). The following is supported:

- Up to 2000 IP addresses for the switch
- Up to 32 IP addresses for the same VLAN
- Up to 512 IP VLANs, that is, VLANs on which you can configure IP addresses
- Each IP address on a VLAN must be for a separate subnet, whether on the same VLAN or different VLANs.

Syntax: [no] vlan < vlan-id > ip address < ip-address/mask-length >
[no] vlan < vlan-id > ip address < ip-address > < mask-bits >

For example, if you wanted to multinet VLAN_20 (VID = 20) with the IP addresses shown below, you would perform steps similar to the following. (For this example, assume that the first IP address is already configured.)

IP Address	VID	IP Address	Subnet Mask
1st address	20	10.25.33.101	255.255.240.0
2nd address	20	10.26.33.101	255.255.240.0
3rd address	20	10.27.33.101	255.255.240.0

1. Go to VLAN 20.
2. Configure two additional IP addresses on VLAN 20.
3. Display IP addressing.

```

ProCurve(config)# vlan 20
ProCurve(vlan-20)# ip address 10.26.33.101/20
ProCurve(vlan-20)# ip address 10.27.33.101/20

ProCurve(vlan-20)# show ip

Internet (IP) Service
IP Routing : Disabled

Default Gateway :
Default TTL    : 64
Arp Age       : 20

VLAN          | IP Config | IP Address   | Subnet Mask
-----+-----
DEFAULT_VLAN | Manual   | 10.20.30.100 | 255.255.240.0
VLAN_20      | Manual   | 10.25.33.101 | 255.255.240.0
              | Manual   | 10.26.33.101 | 255.255.240.0
              | Manual   | 10.27.33.101 | 255.255.240.0

```

Figure 8-4. Example of Configuring and Displaying a Multinetted VLAN

If you then wanted to multinet the default VLAN, you would do the following:

```
ProCurve(vlan-20)# vlan 1
ProCurve(vlan-1)# ip address 10.21.30.100/20
ProCurve(vlan-1)# show ip
Internet (IP) Service
  IP Routing : Disabled
  Default Gateway :
  Default TTL   : 64
  Arp Age      : 20
```

VLAN	IP Config	IP Address	Subnet Mask
DEFAULT_VLAN	Manual	10.20.30.100	255.255.240.0
	Manual	10.21.30.100	
VLAN_20	Manual	10.25.33.101	255.255.240.0
	Manual	10.26.33.101	
	Manual	10.27.33.101	

Figure 8-5. Example of Multinetting on the Default VLAN

Note

The Internet (IP) Service screen in the Menu interface (figure 8-1 on page 8-6) displays the first IP address for each VLAN. You must use the CLI **show ip** command to display the full IP address listing for multinetted VLANs.

Removing or Replacing IP Addresses in a Multinetted VLAN. To remove an IP address from a multinetted VLAN, use the **no** form of the IP address command shown on page 8-9. Generally, to replace one IP address with another, you should first remove the address you want to replace, and then enter the new address.

Configure the Optional Default Gateway. Using the Global configuration level, you can manually assign one default gateway to the switch. (The switch does *not* allow IP addressing received from a DHCP or Bootp server to replace a manually configured default gateway.)

Syntax: ip default-gateway <ip-address >

For example:

```
ProCurve(config)# ip default-gateway 10.28.227.115
```

Note

The switch uses the IP default gateway only while operating as a Layer 2 device. While routing is enabled on the switch, the IP default gateway is not used. Thus, to avoid loss of Telnet access to off-subnet management stations, you should use the **ip route** command to configure a static (default) route before enabling routing. For more information, refer to the chapter titled “IP Routing Features” in the *Multicast and Routing Guide* for your switch.

Configure Time-To-Live (TTL). The maximum number of routers (hops) through which a packet can pass before being discarded. (The default is 64.) Each router decreases a packet’s TTL by 1 before forwarding the packet. If a router decreases the TTL to 0, the router drops the packet instead of forwarding it.

Syntax: ip ttl <number-of-hops>

```
ProCurve(config)# ip ttl 60
```

In the CLI, you can execute this command only from the global configuration level. The TTL default is 64, and the range is 2 - 255.

Web: Configuring IP Addressing

You can use the web browser interface to access IP addressing only if the switch already has an IP address that is reachable through your network.

1. Click on the Configuration tab.
2. Click on **[IP Configuration]**.
3. If you need further information on using the web browser interface, click on **[?]** to access the web-based help available for the switch.

How IP Addressing Affects Switch Operation

Without an IP address and subnet mask compatible with your network, the switch can be managed only through a direct terminal device connection to the Console RS-232 port. You can use direct-connect console access to take advantage of features that do not depend on IP addressing. However, to realize the full capabilities ProCurve proactive networking offers through the switch, configure the switch with an IP address and subnet mask compatible with your network. The following table lists the general features available with and without a network-compatible IP address configured.

Table 8-1. Features Available With and Without IP Addressing on the Switch

Features Available Without an IP Address	Additional Features Available with an IP Address and Subnet Mask
<ul style="list-style-type: none">• Direct-connect access to the CLI and the menu interface.• DHCP or Bootp support for automatic IP address configuration, and DHCP support for automatic Timep server IP address configuration• Multiple Spanning Tree Protocol• Port settings and port trunking• Switch meshing• Console-based status and counters information for monitoring switch operation and diagnosing problems through the CLI or menu interface.• VLANs and GVRP• Serial downloads of software updates and configuration files (Xmodem)• Link test• Port monitoring• Password authentication• Quality of Service (QoS)• Authorized IP manager security	<ul style="list-style-type: none">• Web browser interface access, with configuration, security, and diagnostic tools, plus the Alert Log for discovering problems detected in the switch along with suggested solutions• SNMP network management access such as ProCurve Manager for network configuration, monitoring, problem-finding and reporting, analysis, and recommendations for changes to increase control and uptime• TACACS+, RADIUS, SSH, SSL, and 802.1X authentication• Multinetting on VLANs• Telnet access to the CLI or the menu interface• IGMP• TimeP and SNTP server configuration• TFTP download of configurations and software updates• Access Control Lists (ACLs)• IP routing, Multicast Routing• VRRP router redundancy• PIM-DM and PIM-SM• Radius• Ping test

DHCP/Bootp Operation

Overview. DHCP/Bootp is used to provide configuration data from a DHCP or Bootp server to the switch. This data can be the IP address, subnet mask, default gateway, Timep Server address, and TFTP server address. If a TFTP server address is provided, this allows the switch to TFTP a previously saved configuration file from the TFTP server to the switch. With either DHCP or Bootp, the servers must be configured prior to the switch being connected to the network.

Note

The switches covered in this guide are compatible with both DHCP and Bootp servers.

The DHCP/Bootp Process. Whenever the **IP Config** parameter in the switch or in an individual VLAN in the switch is configured to **DHCP/Bootp** (the default), or when the switch is rebooted with this configuration:

1. DHCP/Bootp requests are automatically broadcast on the local network. (The switch sends one type of request to which either a DHCP or Bootp server can respond.)
2. When a DHCP or Bootp server receives the request, it replies with a previously configured IP address and subnet mask for the switch. The switch also receives an IP Gateway address if the server has been configured to provide one. In the case of Bootp, the server must first be configured with an entry that has the switch's MAC address. (To determine the switch's MAC address, refer to Appendix D, "MAC Address Management".) The switch properly handles replies from either type of server. If multiple replies are returned, the switch tries to use the first reply.)

Note

If you manually configure default gateway, TTL, TimeP, and/or SNTP parameters on the switch, it ignores any values received for the same parameters via DHCP or Bootp.

If the switch is initially configured for DHCP/Bootp operation (the default), or if it reboots with this configuration, it begins sending request packets on the network. If the switch does not receive a reply to its DHCP/Bootp requests, it continues to periodically send request packets, but with decreasing frequency. Thus, if a DHCP or Bootp server is not available or accessible to the switch when DHCP/Bootp is first configured, the switch may not immediately receive the desired configuration. After verifying that the server has become accessible to the switch, reboot the switch to re-start the process immediately.

DHCP Operation. A significant difference between a DHCP configuration and a Bootp configuration is that an IP address assignment from a DHCP server is automatic. Depending on how the DHCP server is configured, the switch may receive an IP address that is temporarily leased. Periodically the switch may be required to renew its lease of the IP configuration. Thus, the IP addressing provided by the server may be different each time the switch reboots or renews its configuration from the server. However, you can fix the address assignment for the switch by doing either of the following:

- Configure the server to issue an "infinite" lease.
- Using the switch's MAC address as an identifier, configure the server with a "Reservation" so that it will always assign the same IP address to the switch. (For MAC address information, refer to Appendix D, "MAC Address Management".)

For more information on either of these procedures, refer to the documentation provided with the DHCP server.

Bootp Operation. When a Bootp server receives a request it searches its Bootp database for a record entry that matches the MAC address in the Bootp request from the switch. If a match is found, the configuration data in the associated database record is returned to the switch. For many Unix systems, the Bootp database is contained in the **/etc/bootptab** file. In contrast to DHCP operation, Bootp configurations are always the same for a specific receiving device. That is, the Bootp server replies to a request with a configuration previously stored in the server and designated for the requesting device.

Bootp Database Record Entries. A minimal entry in the Bootp table file **/etc/bootptab** to update an IP address and subnet mask to the switch or a VLAN configured in the switch would be similar to this entry:

```
8212switch:\
  ht=ether:\
  ha=0030c1123456:\
  ip=10.66.77.88:\
  sm=255.255.248.0:\
  gw=10.66.77.1:\
  hn:\
  vm=rfc1048
```

An entry in the Bootp table file **/etc/bootptab** to tell the switch or VLAN where to obtain a configuration file download would be similar to this entry:

```
8212switch:\
  ht=ether:\
  ha=0030c1123456:\
  ip=10.66.77.88:\
  sm=255.255.248.0:\
  gw=10.66.77.1:\
  lg=10.22.33.44:\
  T144="switch.cfg":\
  vm=rfc1048
```

where:

8212switch	is a user-defined symbolic name to help you find the correct section of the bootptab file. If you have multiple switches that will be using Bootp to get their IP configuration, you should use a unique symbolic name for each switch.
ht	is the "hardware type". For the switches covered in this guide, enter ether (for Ethernet). <i>This tag must precede the ha tag.</i>
ha	is the "hardware address". Use the switch's (or VLAN's) 12-digit MAC address.
ip	is the IP address to be assigned to the switch (or VLAN).
sm	is the subnet mask of the subnet in which the switch (or VLAN) is installed.
gw	is the IP address of the default gateway.

lg	TFTP server address (source of final configuration file)
T144	is the vendor-specific "tag" identifying the configuration file to download.
vm	is a required entry that specifies the Bootp report format. Use rfc1048 for the switches covered in this guide.

Note

The above Bootp table entry is a sample that will work for the switch when the appropriate addresses and file names are used.

Network Preparations for Configuring DHCP/Bootp

In its default configuration, the switch is configured for DHCP/Bootp operation. However, the DHCP/Bootp feature will not acquire IP addressing for the switch unless the following tasks have already been completed:

- For Bootp operation:
 - A Bootp database record has already been entered into an appropriate Bootp server.
 - The necessary network connections are in place
 - The Bootp server is accessible from the switch
- For DHCP operation:
 - A DHCP scope has been configured on the appropriate DHCP server.
 - The necessary network connections are in place
 - A DHCP server is accessible from the switch

Note

Designating a primary VLAN other than the default VLAN affects the switch's use of information received via DHCP/Bootp. For more on this topic, refer to the chapter describing VLANs in the *Advanced Traffic Management Guide* for your switch.

After you reconfigure or reboot the switch with DHCP/Bootp enabled in a network providing DHCP/Bootp service, the switch does the following:

- Receives an IP address and subnet mask and, if configured in the server, a gateway IP address and the address of a Timep server.
- If the DHCP/Bootp reply provides information for downloading a configuration file, the switch uses TFTP to download the file from the designated source, then reboots itself. (This assumes that the switch or VLAN has connectivity to the TFTP file server specified in the reply, that the configuration file is correctly named, and that the configuration file exists in the TFTP directory.)

Loopback Interfaces

This section describes how to configure and use user-defined loopback interfaces on the switch.

Introduction

By default, each switch has an internal loopback interface (**lo0**) with the IP address 127.0.0.1. This IP address is used only for internal traffic transmitted within the switch and is not used in packet headers in egress traffic sent to network devices.

You can configure up to seven other loopback interfaces (**lo1**, **lo2**, **lo3**, and so on) on the switch to use to transmit network across the network. Each loopback interface can have multiple IP addresses. Routing protocols, such as RIP and OSPF, advertise the configured loopback addresses throughout a network or autonomous system.

User-defined loopback addresses provide the following benefits:

- A loopback interface is a virtual interface that is always up and reachable as long as at least one of the IP interfaces on the switch is operational. As a result, a loopback interface is useful for debugging tasks since its IP address can always be pinged if any other switch interface is up.
- You can use a loopback interface to establish a Telnet session, ping the switch, and access the switch through SNMP, SSH, and HTTP (web interface).
- A loopback IP address can be used by routing protocols. For example, you can configure the loopback IP address as the router ID used to identify the switch in an OSPF area. Because the loopback interface is always up, you ensure that the switch's router ID remains constant and that the OSPF network is protected from changes caused by downed interfaces.

Note

OSPF does not require that you use an IP address as the router ID. OSPF only requires the router ID to be a unique value within the autonomous system (AS). However, if you configure the loopback IP address as the router ID, OSPF can reach the switch if any switch interface is up. (Normally, OSPF automatically configures the router ID with the IP address of a switch interface. The disadvantage is that if the interface goes down, OSPF can no longer ping the switch using the router ID even if other interfaces are operational.)

For more information about how to configure a loopback IP address to participate in an OSPF broadcast area, refer to the section titled “(Optional) Assigning Loopback Addresses to an Area” in the *Multicast and Routing Guide*.

Configuring a Loopback Interface

To configure a loopback interface, enter the **interface loopback** command at the global configuration level of the CLI:

Syntax: [no] interface loopback <number>

*Creates a loopback interface, where <number> is a value from 1 to 7. Use the **no** form of the command to remove the loopback interface.*

Note: You cannot remove the default loopback interface (number 0) with IP address 127.0.0.1.

You can configure up to thirty-two IP addresses on a loopback interface. To configure an IP address for the loopback interface, enter the **ip address <ip-address>** command at the loopback interface configuration level as shown in the following example.

Note that when you configure an IP address for a loopback interface, you do not specify a network mask. The default subnet mask 255.255.255.255 is used.

```
ProCurve(config)# interface loopback 1
ProCurve (lol)# ip address 10.1.1.1
```

Figure 8-6. Example of a Loopback Interface Configuration

Notes

- You can configure a loopback interface only from the CLI; you cannot configure a loopback interface from the web management or Menu interface.
- Loopback interfaces share the same IP address space with VLAN configurations. The maximum number of IP addresses supported on a switch is 2048, which includes all IP addresses configured for both VLANs and loopback interfaces (except for the default loopback IP address 127.0.0.1).
- Each IP address that you configure on a loopback interface must be unique in the switch. This means that the address cannot be used by a VLAN interface or another loopback interface.

For example, if you configure a VLAN with IP address 172.16.100.8/24, you cannot configure a loopback interface with IP address 172.16.100.8. In the same way, if you configure a loopback interface (**lo1**) with IP address 172.16.101.8, you cannot configure another loopback interface (**lo2**) with IP address 172.16.101.8.

- You can configure multiple IP addresses on a loopback interface (**lo0** to **lo7**). Up to thirty-two IP addresses are supported on a loopback interface. The following example shows valid IP address configurations on two loopback interfaces.

```
ProCurve(config)# interface loopback 0
ProCurve (lo0)# ip address 172.16.101.8
ProCurve (lo0)# ip address 172.16.101.9
ProCurve (lo0)# exit
ProCurve (config)# interface loopback 1
ProCurve (lo1)# ip address 172.16.102.1
ProCurve (lo1)# ip address 172.16.102.2
```

Displaying Loopback Interface Configurations

To display the list of loopback interfaces which have been assigned IP addresses, enter the **show ip** command.

In the **show ip** command output, information about configured loopback interfaces is displayed below other IP configuration parameters, such as packet time-to-live (TTL) and ARP age-out values, and VLAN IP configurations. The following example displays the IP addresses configured for two user-defined loopback interfaces (**lo1** and **lo2**).


```

ProCurve> show ip

Internet (IP) Service

IP Routing : Enabled
Default TTL : 64
ARP Age : 20

VLAN          IP Config IP Address Subnet Mask Proxy ARP
-----
DEFAULT_VLAN  Manual  10.0.8.121  255.255.0.0  No
VLAN2         Manual  192.168.12.1 255.255.255.0 No
VLAN3         Disabled

Loopback      Loopback Addresses
Loopback      IP Config  IP Address  Subnet Mask
-----
lo1           Manual     172.16.110.2 255.255.255.255
lo2           Manual     172.16.112.2 255.255.255.255
lo2           Manual     172.16.114.1 255.255.255.255

```

Figure 8-7. Example of show ip Command Output

Note

The default loopback interface (**lo0**) with IP address 127.0.0.1 is not displayed in the **show ip** command output because it is permanently configured on the switch. To display the default loopback address, enter the **show ip route** command as shown in figure 8-8.

To display the loopback interfaces configured on the switch in a list of IP routing entries displayed according to destination IP address, enter the **show ip route** command.

The following example displays the configuration of the default loopback interface (**lo0**) and one user-defined loopback interface (**lo2**).

```
ProCurve> show ip route

IP Route Entries

IP Routing : Enabled
Default TTL : 64
ARP Age : 20

Destination      Gateway          VLAN Type      Metric  Dist
-----
10.0.0.0/16      DEFAULT_VLAN    1   connected  1       0
127.0.0.0/8      reject          static  0       0
127.0.0.1/32     lo0             connected  1       0
172.16.10.121/32 lo2             static  1       0
172.16.100.0/24  10.0.8.11      1   ospf       1       1
172.16.102.0/24  VLAN2          2   connected  1       0
```

Figure 8-8. Example of show ip route Command Output

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

For the switches covered in this guide, IP Preserve enables you to copy a configuration file to multiple switches while retaining the individual IP address and subnet mask on VLAN 1 in each switch, and the Gateway IP address assigned to the switch. This enables you to distribute the same configuration file to multiple switches without overwriting their individual IP addresses.

Operating Rules for IP Preserve

When **ip preserve** is entered as the last line in a configuration file stored on a TFTP server:

- If the switch's current IP address for VLAN 1 was not configured by DHCP/Bootp, IP Preserve retains the switch's current IP address, subnet mask, and IP gateway address when the switch downloads the file and reboots. The switch adopts all other configuration parameters in the configuration file into the startup-config file.
- If the switch's current IP addressing for VLAN 1 is from a DHCP server, IP Preserve is suspended. In this case, whatever IP addressing the configuration file specifies is implemented when the switch downloads the file and reboots. If the file includes DHCP/Bootp as the IP addressing source for VLAN 1, the switch will configure itself accordingly and use DHCP/Bootp. If instead, the file includes a dedicated IP address and subnet mask for VLAN 1 and a specific gateway IP address, then the switch will implement these settings in the startup-config file.
- The **ip preserve** statement does not appear in **show config** listings. To verify IP Preserve in a configuration file, open the file in a text editor and view the last line. For an example of implementing IP Preserve in a configuration file, see figure 8-9, below.

Enabling IP Preserve

To set up IP Preserve, enter the **ip preserve** statement at the end of a configuration file. (Note that you do not execute IP Preserve by entering a command from the CLI).

```
; J8697A Configuration Editor; Created on release #K.12.00
hostname "ProCurve"
time daylight-time-rule None
-
.
.
.
password manager
password operator
ip preserve
```

Entering "ip preserve" in the last line of a configuration file implements IP Preserve when the file is downloaded to the switch and the switch reboots.

Figure 8-9. Example of Implementing IP Preserve in a Configuration File

For example, consider figure 8-10:

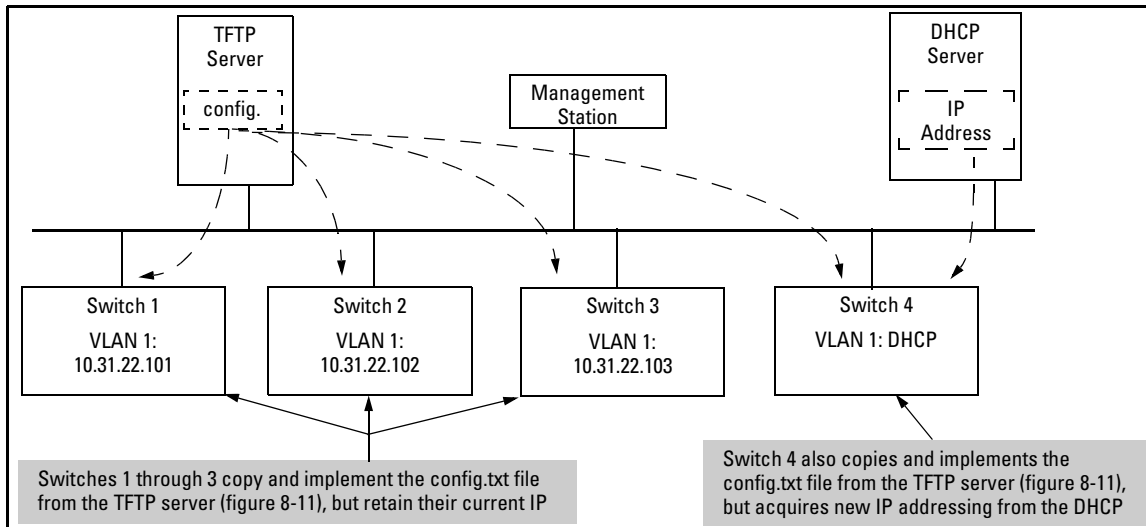


Figure 8-10. Example of IP Preserve Operation with Multiple Series Switches

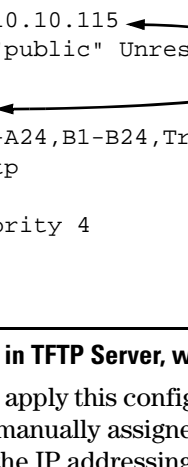
If you apply the following configuration file to figure 8-10, switches 1 - 3 will retain their manually assigned IP addressing and switch 4 will be configured to acquire its IP addressing from a DHCP server.

```
ProCurve(config)# show run

Running configuration:

; J8715A Configuration Editor; Created on release #K.12.07

hostname "ProCurve"
module 1 type J8702A
module 2 type J8705A
trunk A11-A12 Trk1 Trunk
ip default-gateway 10.10.10.115
snmp-server community "public" Unrestricted
vlan 1
  name "DEFAULT_VLAN"
  untagged A1-A10,A13-A24,B1-B24,Trk1
  ip address dhcp-bootp
  exit
spanning-tree Trk1 priority 4
password manager
password operator
```



Using figure 8-10, above, switches 1 - 3 ignore these entries because the file implements IP Preserve and their current IP addressing was not acquired through DHCP/Bootp.

Switch 4 ignores IP Preserve and implements the DHCP/Bootp addressing and IP Gateway specified in this file (because its last IP addressing was acquired from a DHCP/Bootp server).

Figure 8-11. Configuration File in TFTP Server, with DHCP/Bootp Specified as the IP Addressing Source

If you apply this configuration file to figure 8-10, switches 1 - 3 will still retain their manually assigned IP addressing. However, switch 4 will be configured with the IP addressing included in the file.

Configuring IP Addressing

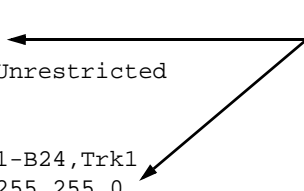
IP Preserve: Retaining VLAN-1 IP Addressing Across Configuration File Downloads

```
ProCurve# show run

Running configuration:

; J8715A Configuration Editor; Created on release #K.12.07

hostname "ProCurve"
module 1 type J8702A
module 2 type J8705A
trunk A11-A12 Trk1 Trunk
ip default-gateway 10.10.10.115
snmp-server community "public" Unrestricted
vlan 1
  name "DEFAULT_VLAN"
  untagged A1,A7-A10,A13-A24,B1-B24,Trk1
  ip address 10.12.17.175 255.255.255.0
  tagged A4-A6
  no untagged A2-A3
  exit
vlan 2
  name "VLAN2"
  untagged A2-A3
  no ip address
  exit
spanning-tree Trk1 priority 4
password manager
password operator
```



Because switch 4 (figure 8-10) received its most recent IP addressing from a DHCP/Bootp server, the switch ignores the **ip preserve** command and implements the IP addressing included in this file.

Figure 8-12. Configuration File in TFTP Server, with Dedicated IP Addressing Instead of DHCP/Bootp

To summarize the IP Preserve effect on IP addressing:

- If the switch received its most recent VLAN 1 IP addressing from a DHCP/Bootp server, it ignores the IP Preserve command when it downloads the configuration file, and implements whatever IP addressing instructions are in the configuration file.
- If the switch did not receive its most recent VLAN 1 IP addressing from a DHCP/Bootp server, it retains its current IP addressing when it downloads the configuration file.
- The content of the downloaded configuration file determines the IP addresses and subnet masks for other VLANs.

Configuring a Single Source IP Address

Overview

This feature applies to the following software applications:

- TACACS
- RADIUS
- System Logging applications

The above IP-based software applications use a client-server communication model, that is, the client's source IP address is used for unique client identification. The source IP address is determined by the system and is usually the IP address of the outgoing interface in the routing table. However, routing switches may have multiple routing interfaces due to load balancing or routing redundancy, and outgoing packets can potentially be sent by different paths at different times. This results in different source IP addresses, which creates a client identification problem on the server site. For example, there is no way to designate a fixed IP address for outgoing packets for RADIUS or TACACS, so it is necessary to configure in the RADIUS or TACACS database all possible IP addresses that are configured on the switch as valid clients. When using system logging, it can be difficult to interpret the logging and accounting data on the server site as the same client can be logged with different IP addresses.

To decrease the amount of administrative work involved, a configuration model is provided that allows the selection of an IP address to use as the source address for all outgoing traffic generated by a specified software application on the switch. This allows unique identification of the software application on the server site regardless of which local interface has been used to reach the destination server.

Specifying the Source IP Address

The CLI command **ip source-interface** is used to specify the source IP address for an application. Different source IP addresses can be used for different software applications, but only one source IP address can be specified for each application.

Syntax: [no] ip source-interface <radius | tacacs | logging | all> <loopback <id> | vlan <vlan-id> address <ip-address>>

*Determines the source IP address used by the specified software application when transmitting IP packets. The **all** parameter can be used to set one IP address for all the listed applications, in this case, RADIUS, TACACS, and System Logging.*

*The **no** version of the command cancels the configuration and the application reverts to its default behavior. The system determines the source IP address of outgoing application-specific IP packets at packet transmission time.*

loopback <id>: *Specifies that the IP address of the loopback interface is used as the source IP address in outgoing packets. If the loopback interface has no IP address, then the application reverts to the default behavior. If more than one IP address is configured, then the lowest IP address is used.*

vlan <vlan-id>: *Specifies that the IP address of the indicated VLAN interface is used as the source IP address of outgoing packets. If the specified VLAN interface has no IP address configured, or is down, then the application reverts to the default behavior. If more than one IP address is configured, then the lowest IP address is used.*

address <ip-address>: *Specifies the IP address that should be used as the source IP address of outgoing packets. The IP address must be a valid IP address configured on one of the switch's VLAN or loopback interfaces. If the interface is down, then the application reverts to the default behavior.*

The Source IP Selection Policy

The source IP address selection for the application protocols is defined through assignment of one of the following policies:

- **Outgoing Interface**—the IP address of the outgoing IP interface is used as the source IP address. This is the default policy and the default behavior of applications.

- Configured IP Address—the specific IP address that is used as the source IP address. This address is configured on one of the switch’s IP interfaces, either a VLAN interface or a Loopback interface.
- Configured IP Interface—the IP address from the specific IP interface (VLAN or Loopback) is used as the source IP address. If there are multiple IP addresses assigned (multinetting, for example), the lowest IP address is used.

If the selection policy cannot be executed because the interface does not have an IP address configured, does not exist, or is down, the application protocol uses the default Outgoing Interface policy. A warning message is displayed, but the configuration changes are accepted. When using the **show ip source-interface status** command to display information about the source IP address selection policy, the administratively-assigned source IP selection policy and the actual (operational) source IP selection policy in effect are displayed. The operational source IP selection policy may be different from the assigned source selection policy if the IP interface does not exist or is down. In this case, the default of Outgoing Interface appears as the operational policy (See figure 8-13).

```
ProCurve (config)# show ip source-interface detail

Source-IP Detailed Information

Protocol : Tacacs
Admin Policy      : Configured IP Interface
Oper Policy      : Outgoing Interface
Source IP Interface : Vlan 22
Source IP Address  : 10.10.10.4
Source Interface State : Down
```

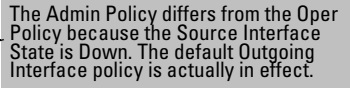


Figure 8-13. Example of the Administratively-assigned Source IP Selection Policy Differing From the Operational Policy

The **no** form of the **ip source-interface** command reverts the application protocols to the default behavior. The Outgoing Interface policy is used.

Figure 8-14 is an example of assigning a specific source IP address for a RADIUS application. The administrative policy is Configured IP Address.

Configuring IP Addressing

Configuring a Single Source IP Address

```
ProCurve(config)# ip source-interface radius address 10.10.10.2
ProCurve(config)# show ip source-interface radius

Source-IP Configuration Information

Protocol | Admin Selection Policy | IP Interface | IP Address
-----+-----
Radius   | Configured IP Address  | vlan 3      | 10.10.10.2
```

Figure 8-14. Example of a Specific IP Address Assigned for the RADIUS Application Protocol

In figure 8-15, a VLAN interface (VLAN 22) is specified as the source IP address for TACACS. The administrative policy is Configured IP Interface.

```
ProCurve(config)# ip source-interface tacacs vlan 22
ProCurve(config)# show ip source-interface tacacs

Source-IP Configuration Information

Protocol | Admin Selection Policy | IP Interface | IP Address
-----+-----
Tacacs   | Configured IP Interface | vlan 22     | 10.10.10.4
```

Figure 8-15. Example of Using a VLAN Interface as the Source IP Address for TACACS

Figure 8-16 shows a VLAN interface being specified as the source IP address for logging. The administrative policy is Configured IP Interface.

```
ProCurve(config)# ip source-interface syslog vlan 10
ProCurve(config)# show ip source-interface syslog

Source-IP Configuration Information

Protocol | Admin Selection Policy | IP Interface | IP Address
-----+-----
Syslog   | Configured IP Interface | vlan 10     | 10.10.10.10
```

Figure 8-16. Example of Using a VLAN Interface as the Source IP Address for Logging (Syslog)

Displaying the Source IP Interface Information

There are several **show** commands that can be used to display information about the source IP interface status.

Syntax: show ip source-interface status [radius | tacacs | syslog]

Displays the operational status information for the source IP address selection policy. Both the administratively-assigned source IP selection policy and the operational source IP selection policy are displayed.

When no parameters are specified, policy information for all protocols is displayed.

```
ProCurve(config)# show ip source-interface status

Source-IP Status Information

Protocol | Admin Selection Policy  Oper Selection Policy
-----+-----
Tacacs   | Configured IP Interface Configured IP Interface
Radius   | Configured IP Address   Configured IP Address
Syslog   | Configured IP Interface Outgoing Interface
```

Figure 8-17. Example of the Data Displayed for Source IP Interface Status

When executing the **show ip source-interface** command without parameters, the configured IP interfaces (VLANs) and IP addresses are displayed for each protocol.

```
ProCurve(config)# show ip source-interface

Source-IP Configuration Information

Protocol | Admin Selection Policy  IP Interface  IP Address
-----+-----
Tacacs   | Configured IP Interface vlan 22      10.10.10.4
Radius   | Configured IP Address   vlan 3        10.10.10.2
Syslog   | Configured IP Interface vlan 10      10.10.10.10
```

Figure 8-18. Example of show ip source-interface Command Output

The **show ip source-interface detail** command displays detailed information about the configured policies, source IP address, and interface state for each protocol.

Syntax: show ip source-interface detail [radius | tacacs | syslog]

Displays detailed operational status information for the source IP address selection policy. Information about the configured policies, source IP address and interface state are displayed.

When no parameters are specified, policy information for all protocols is displayed.

```
ProCurve(config)# show ip source-interface detail

Source-IP Detailed Information

Protocol : Tacacs
Admin Policy           : Configured IP Interface
Oper Policy            : Configured IP Interface
Source IP Interface    : vlan 22
Source IP Address      : 10.10.10.4
Source Interface State : Up

Protocol : Radius
Admin Policy           : Configured IP Address
Oper Policy            : Configured IP Address
Source IP Interface    : vlan 3
Source IP Address      : 10.10.10.2
Source Interface State : Up

Protocol : Syslog
Admin Policy           : Configured IP Interface
Oper Policy            : Configured IP Interface
Source IP Interface    : vlan 10
Source IP Address      : 10.10.10.10
Source Interface State : Up
```

Figure 8-19. Example of Detailed Information Displayed for Each Protocol

The **show** command can also be used with the application to display the source IP address selection information in effect for the application protocol.

```
ProCurve(config)# show radius

Status and Counters - General RADIUS Information

Deadttime(min) : 0
Timeout(secs) : 5
Retransmit Attempts : 3
Global Encryption Key :
Dynamic Authorization UDP Port : 3799
Source IP Selection : Configured IP address
```

← Source IP Selection for the specified application protocol is displayed.

Figure 8-20. Example of show radius Command Displaying Source IP Selection Information

```
ProCurve(config)# show tacacs

Status and Counters - TACACS Information

Timeout : 5
Source IP Selection : Configured IP Interface
Encryption Key :
```

← Source IP Selection for the specified application protocol is displayed.

Figure 8-21. Example of show tacacs Command Displaying Source IP Selection Information

```
ProCurve(config)# show debug

Debug Logging

Source IP Selection: Configured IP interface
Destination:      None

Enabled debug types:
None are enabled.
```

← Source IP Selection for the specified application protocol is displayed.

Figure 8-22. Example of show debug Command Displaying Source IP Selection Information for Syslog

Error Messages

The following error messages may appear when configuring source IP selection if the interface does not exist, is not configured for IP, or is down.

Error Message	Description
Warning: Specified IP address is not configured on any interface	The IP address specified has not been assigned to any interface on the switch.
Warning: Specified IP interface is not configured	The IP interface has not been configured.
Warning: Specified IP interface is not configured for IP	An IP address has not been assigned to this interface.
Warning: Specified IP interface is down.	The interface on the switch associated with this IP address is down.
Warning: Specified IP interface is configured for DHCP	The IP address has not been configured specifically (manually) for this interface and may change.

Time Protocols

Contents

Overview	9-2
TimeP Time Synchronization	9-2
SNTP Time Synchronization	9-2
Selecting a Time Synchronization Protocol or Turning Off Time Protocol Operation	9-3
General Steps for Running a Time Protocol on the Switch:	9-3
Disabling Time Synchronization	9-4
SNTP: Viewing, Selecting, and Configuring	9-4
Menu: Viewing and Configuring SNTP	9-5
CLI: Viewing and Configuring SNTP	9-8
Viewing the Current SNTP Configuration	9-8
Configuring (Enabling or Disabling) the SNTP Mode	9-10
SNTP Client Authentication	9-15
Requirements	9-16
Configuring the Key-Identifier, Authentication Mode, and Key-Value	9-17
Configuring a Trusted Key	9-18
Associating a Key with an SNTP Server	9-19
Enabling SNTP Client Authentication	9-19
Configuring Unicast and Broadcast Mode	9-20
Displaying SNTP Configuration Information	9-20
Saving Configuration Files and the Include-Credentials Command	9-22
.....	9-24
TimeP: Viewing, Selecting, and Configuring	9-25
Menu: Viewing and Configuring TimeP	9-26
CLI: Viewing and Configuring TimeP	9-27

Viewing the Current TimeP Configuration	9-28
Configuring (Enabling or Disabling) the TimeP Mode	9-29
SNTP Unicast Time Polling with Multiple SNTP Servers	9-34
Displaying All SNTP Server Addresses Configured on the Switch ..	9-34
Adding and Deleting SNTP Server Addresses	9-35
Menu: Operation with Multiple SNTP Server Addresses Configured	9-35
SNTP Messages in the Event Log	9-35

Overview

This chapter describes:

- SNTP Time Protocol Operation
- Timep Time Protocol Operation

Using time synchronization ensures a uniform time among interoperating devices. This helps you to manage and troubleshoot switch operation by attaching meaningful time data to event and error messages.

The switch offers TimeP and SNTP (Simple Network Time Protocol) and a **timesync** command for changing the time protocol selection (or turning off time protocol operation).

Notes

- Although you can create and save configurations for both time protocols without conflicts, the switch allows only one active time protocol at any time.
- In the factory-default configuration, the time synchronization option is set to TimeP, with the TimeP mode itself set to **Disabled**.

TimeP Time Synchronization

You can either manually assign the switch to use a TimeP server or use DHCP to assign the TimeP server. In either case, the switch can get its time synchronization updates from only one, designated Timep server. This option enhances security by specifying which time server to use.

SNTP Time Synchronization

SNTP provides two operating modes:

- **Broadcast Mode:** The switch acquires time updates by accepting the time value from the first SNTP time broadcast detected. (In this case, the SNTP server must be configured to broadcast time updates to the network broadcast address. Refer to the documentation provided with your SNTP server application.) Once the switch detects a particular server, it ignores time broadcasts from other SNTP servers unless the configurable **Poll Interval** expires three consecutive times without an update received from the first-detected server.

Note

To use Broadcast mode, the switch and the SNTP server must be in the same subnet.

- **Unicast Mode:** The switch requests a time update from the configured SNTP server. (You can configure one server using the menu interface, or up to three servers using the CLI **sntp server** command.) This option provides increased security over the Broadcast mode by specifying which time server to use instead of using the first one detected through a broadcast.
-

Selecting a Time Synchronization Protocol or Turning Off Time Protocol Operation

General Steps for Running a Time Protocol on the Switch:

1. Select the time synchronization protocol: **SNTP** or **TimeP** (the default).
2. Enable the protocol. The choices are:
 - SNTP: **Broadcast** or **Unicast**
 - TimeP: **DHCP** or **Manual**
3. Configure the remaining parameters for the time protocol you selected.

The switch retains the parameter settings for both time protocols even if you change from one protocol to the other. Thus, if you select a time protocol, the switch uses the parameters you last configured for the selected protocol.

Note that simply selecting a time synchronization protocol does not enable that protocol on the switch unless you also enable the protocol itself (step 2, above). For example, in the factory-default configuration, TimeP is the selected time synchronization method. However, because TimeP is disabled in the factory-default configuration, no time synchronization protocol is running.

Disabling Time Synchronization

You can use either of the following methods to disable time synchronization without changing the Timep or SNTP configuration:

- In the System Information screen of the Menu interface, set the **Time Synch Method** parameter to **None**, then press **[Enter]**, then **[S]** (for **Save**).
- In the Global config level of the CLI, execute **no timesync**.

SNTP: Viewing, Selecting, and Configuring

SNTP Feature	Default	Menu	CLI	Web
view the SNTP time synchronization configuration	n/a	page 9-6	page 9-9	—
select SNTP as the time synchronization method	timep	page 9-7	page 9-11 ff.	—
disable time synchronization	timep	page 9-7	page 9-15	—
enable the SNTP mode (Broadcast, Unicast, or Disabled)	disabled			—
broadcast	n/a	page 9-7	page 9-12	—
unicast	n/a	page 9-7	page 9-12	—
none/disabled	n/a	page 9-7	page 9-16	—
configure an SNTP server address (for Unicast mode only)	none	page 9-7	page 9-12 ff.	—
change the SNTP server version (for Unicast mode only)	3	page 9-8	page 9-14	—
change the SNTP poll interval	720 seconds	page 9-8	page 9-15	—
change the server priority	n/a	—	page 9-15	—

Time Protocols

SNTP: Viewing, Selecting, and Configuring

Table 9-1. SNTP Parameters

SNTP Parameter	Operation
Time Sync Method	Used to select either SNTP, TIMEP, or None as the time synchronization method.
SNTP Mode	
Disabled	The Default. SNTP does not operate, even if specified by the Menu interface Time Sync Method parameter or the CLI timesync command.
Unicast	Directs the switch to poll a specific server for SNTP time synchronization. Requires at least one server address.
Broadcast	Directs the switch to acquire its time synchronization from data broadcast by any SNTP server to the network broadcast address. The switch uses the first server detected and ignores any others. However, if the Poll Interval expires three times without the switch detecting a time update from the original server, it the switch accepts a broadcast time update from the next server it detects.
Poll Interval (seconds)	In Unicast Mode: Specifies how often the switch polls the designated SNTP server for a time update. In Broadcast Mode: Specifies how often the switch polls the network broadcast address for a time update. Value between 30-720 seconds.
Server Address	Used only when the SNTP Mode is set to Unicast . Specifies the IP address of the SNTP server that the switch accesses for time synchronization updates. You can configure up to three servers; one using the menu or CLI, and two more using the CLI. Refer to “SNTP Unicast Time Polling with Multiple SNTP Servers” on page 9-35.
Server Version	Default: 3; range: 1 - 7. Specifies the SNTP software version to use, and is assigned on a per-server basis. The version setting is backwards-compatible. For example, using version 3 means that the switch accepts versions 1 through 3.
Priority	Specifies the order in which the configured servers are polled for getting the time. Value is between 1 and 3.

Menu: Viewing and Configuring SNTP

To View, Enable, and Modify SNTP Time Protocol:

1. From the Main Menu, select:
 - 2. Switch Configuration...**
 - 1. System Information**

```
===== CONSOLE - MANAGER MODE =====
Switch Configuration - System Information

System Name : ProCurve
System Contact :
System Location :

Inactivity Timeout (min) [0] : 0      MAC Age Time (sec) [300] : 300
Inbound Telnet Enabled [Yes] : Yes    Web Agent Enabled [Yes] : Yes
Time Sync Method [None] : TIMEP
TimeP Mode [Disabled] : Disabled
Tftp-enable [Yes] : Yes
Time Zone [0] : 0
Daylight Time Rule [None] : None

Server Address :
Jumbo Max Frame Size [9216] : 9216
Jumbo IP MTU [9198] : 9198

Time Protocol Selection Parameter
- TIMEP
- SNTP
- None

Actions->  Cancel      Edit      Save      Help
```

Figure 9-1. The System Information Screen (Default Values)

2. Press **[E]** (for **Edit**). The cursor moves to the **System Name** field.
3. Use **[↓]** to move the cursor to the **Time Sync Method** field.
4. Use the Space bar to select **SNTP**, then press **[↓]** once to display and move to the **SNTP Mode** field.
5. Do one of the following:
 - Use the Space bar to select the **Broadcast** mode, then press **[↓]** to move the cursor to the **Poll Interval** field, and go to step 6. (For Broadcast mode details, refer to “SNTP Operating Modes” on page 9-3.)

```
Time Sync Method [None] : SNTP
SNTP Mode [Disabled] : Broadcast
Poll Interval (sec) [720] : 720
Tftp-enable [Yes] : Yes
Time Zone [0] : 0
Daylight Time Rule [None] : None
```

Figure 9-2. Time Configuration Fields for SNTP with Broadcast Mode

- Use the Space bar to select the **Unicast** mode, then do the following:
 - i. Press **[→]** to move the cursor to the **Server Address** field.
 - ii. Enter the IP address of the SNTP server you want the switch to use for time synchronization.

Time Protocols

SNTP: Viewing, Selecting, and Configuring

Note: This step replaces any previously configured server IP address. If you will be using backup SNTP servers (requires use of the CLI), then refer to “SNTP Unicast Time Polling with Multiple SNTP Servers” on page 9-35.

- iii. Press **↓** to move the cursor to the **Server Version** field. Enter the value that matches the SNTP server version running on the device you specified in the preceding step (step ii). If you are unsure which version to use, ProCurve recommends leaving this value at the default setting of **3** and testing SNTP operation to determine whether any change is necessary.

Note: Using the menu to enter the IP address for an SNTP server when the switch already has one or more SNTP servers configured causes the switch to delete the primary SNTP server from the server list and to select a new primary SNTP server from the IP address(es) in the updated list. For more on this topic, refer to “SNTP Unicast Time Polling with Multiple SNTP Servers” on page 9-35.

- iv. Press **→** to move the cursor to the **Poll Interval** field, then go to step 6.

```
Time Sync Method [None] : SNTP
SNTP Mode [Disabled] : Unicast           Server Address : 10.28.227.15
Poll Interval (sec) [720] : 720         Server Version [3] : 3
Tftp-enable [Yes] : Yes
Time Zone [0] : 0
Daylight Time Rule [None] : None
```

Note: The Menu interface lists only the highest priority SNTP server, even if others are configured. To view all SNTP servers configured on the switch, use the CLI **show management** command. Refer to “SNTP Unicast Time Polling with Multiple SNTP Servers” on page 9-35.

Figure 9-3. SNTP Configuration Fields for SNTP Configured with Unicast Mode

6. In the **Poll Interval** field, enter the time in seconds that you want for a Poll Interval. (For Poll Interval operation, see table 9-1, “SNTP Parameters”, on page 9-6.)
7. Press **[Enter]** to return to the Actions line, then **[S]** (for **Save**) to enter the new time protocol configuration in both the startup-config and running-config files.

CLI: Viewing and Configuring SNTP

CLI Commands Described in this Section

SNTP Command	Page
show sntp	9-9
[no] timesync	9-11 and ff., 9-15
sntp broadcast	9-12
sntp unicast	9-12
sntp server	9-12 and ff.
Protocol Version	9-14
Priority	9-15
poll-interval	9-15
no sntp	9-16

This section describes how to use the CLI to view, enable, and configure SNTP parameters.

Viewing the Current SNTP Configuration

Syntax: show sntp

*This command lists both the time synchronization method (**TimeP**, **SNTP**, or **None**) and the SNTP configuration, even if SNTP is not the selected time protocol.*

For example, if you configured the switch with SNTP as the time synchronization method, then enabled SNTP in broadcast mode with the default poll interval, show sntp lists the following:

Time Protocols

SNTP: Viewing, Selecting, and Configuring

```
ProCurve(config)# show sntp

SNTP Configuration

Time Sync Mode: Sntp
SNTP Mode : Unicast
Poll Interval (sec) [720] : 719

Priority SNTP Server Address                                Protocol Version
-----
1         2001:db8::215:60ff:fe79:8980                    7
2         10.255.5.24                                      3
3         fe80::123%vlan10                                3
```

Figure 9-4. Example of SNTP Configuration When SNTP Is the Selected Time Synchronization Method

In the factory-default configuration (where TimeP is the selected time synchronization method), **show sntp** still lists the SNTP configuration even though it is not currently in use. For example:

```
ProCurve(config)# show sntp

SNTP Configuration

Time Sync Mode: TimeP
SNTP Mode : Unicast
Poll Interval (sec) [720] : 719

Priority SNTP Server Address                                Protocol Version
-----
1         2001:db8::215:60ff:fe79:8980                    7
2         10.255.5.24                                      3
3         fe80::123%vlan10                                3
```

Even though, in this example, TimeP is the current time synchronous method, the switch maintains the SNTP configuration.

Figure 9-5. Example of SNTP Configuration When SNTP Is Not the Selected Time Synchronization Method

Syntax: show management

This command can help you to easily examine and compare the IP addressing on the switch. It lists the IP addresses for all time servers configured on the switch, plus the IP addresses and default gateway for all VLANs configured on the switch.


```

ProCurve(config)# show management

Status and Counters - Management Address Information

Time Server Address : fe80::215:60ff:fe7a:adc0%vlan10

Priority  SNTP Server Address                                Protocol Version
-----  -
1         2001:db8::215:60ff:fe79:8980                        7
2         10.255.5.24                                           3
3         fe80::123%vlan10                                3

Default Gateway      : 10.0.9.80

VLAN Name    MAC Address      | IP Address
-----  -
DEFAULT_VLAN 001279-88a100   | Disabled
VLAN10      001279-88a100   | 10.0.10.17

```

Figure 9-6. Example of Display Showing IP Addressing for All Configured Time Servers and VLANs

Configuring (Enabling or Disabling) the SNTP Mode

Enabling the SNTP mode means to configure it for either broadcast or unicast mode. Remember that to run SNTP as the switch's time synchronization protocol, you must also select SNTP as the time synchronization method by using the CLI **timesync** command (or the Menu interface **Time Sync Method** parameter).

Syntax: `timesync sntp`
Selects SNTP as the time protocol.

sntp < broadcast | unicast >
Enables the SNTP mode (below and page 9-12).

Syntax: `sntp server < ip-addr >`
Required only for unicast mode page 9-12).

Syntax: `sntp server priority <1 - 3 >`
Specifies the order in which the configured servers are polled for getting the time. Value is between 1 and 3.

Syntax: `sntp poll-interval < 30 - 720 >`
Enabling the SNTP mode also enables the SNTP poll interval (default: 720 seconds; page 9-15).

Enabling SNTP in Broadcast Mode. Because the switch provides an SNTP polling interval (default: 720 seconds), you need only these two commands for minimal SNTP broadcast configuration:

Syntax: `timesync sntp`

Selects SNTP as the time synchronization method.

Syntax: `sntp broadcast`

*Configures **broadcast** as the SNTP mode.*

For example, suppose:

- Time synchronization is in the factory-default configuration (TimeP is the currently selected time synchronization method).
- You want to:
 1. View the current time synchronization.
 2. Select SNTP as the time synchronization mode.
 3. Enable SNTP for Broadcast mode.
 4. View the SNTP configuration again to verify the configuration.

The commands and output would appear as follows:

```
ProCurve(config)# show sntp 1
SNTP Configuration
  Time Sync Mode: Timep
  SNTP Mode : disabled
  Poll Interval (sec) [720] : 720
ProCurve(config)# timesync sntp 2
ProCurve(config)# sntp broadcast 3
ProCurve(config)# show sntp 4
```

1 `show sntp` displays the SNTP configuration and also shows that TimeP is the currently active time synchronization mode.

4 `show sntp` again displays the SNTP configuration and shows that SNTP is now the currently active time synchronization mode and is configured for broadcast operation.

Figure 9-7. Example of Enabling SNTP Operation in Broadcast Mode

Enabling SNTP in Unicast Mode. Like broadcast mode, configuring SNTP for unicast mode enables SNTP. However, for Unicast operation, you must also specify the IP address of at least one SNTP server. The switch allows up to three unicast servers. You can use the Menu interface or the CLI to configure one server or to replace an existing Unicast server with another. To add a

second or third server, you must use the CLI. For more on SNTP operation with multiple servers, refer to “SNTP Unicast Time Polling with Multiple SNTP Servers” on page 9-35.

Syntax: `timesync sntp`

Selects SNTP as the time synchronization method.

`sntp unicast`

Configures the SNTP mode for Unicast operation.

Syntax: `[no] sntp server priority <1-3> <ip-address> [oobm] [version]`

*Use the **no** version of the command to disable SNTP.*

priority specifies the order in which the configured SNTP servers are polled for the time; allowable values are 1 through 3.

ip-address is an IPv4 or IPv6 address of an SNTP server.

*For switches that have a separate out-of-band management port, **oobm** specifies that SNTP traffic goes through that port. (By default, SNTP traffic goes through the data ports.)*

version is the protocol version of the SNTP server. Allowable values are 1 through 7; default is 3.

Syntax: `no sntp server < ip-addr >`

Deletes the specified SNTP server.

Note

Deleting an SNTP server when only one is configured disables SNTP unicast operation.

For example, to select SNTP and configure it with unicast mode and an SNTP server at 10.28.227.141 with the default server version (3) and default poll interval (720 seconds):

```
ProCurve(config)# timesync sntp
```

Selects SNTP.

```
ProCurve(config)# sntp unicast
```

Activates SNTP in Unicast mode.

Time Protocols

SNTP: Viewing, Selecting, and Configuring

```
ProCurve(config)# sntp server 10.28.227.141
```

Specifies the SNTP server and accepts the current SNTP server version (default: 3).

```
ProCurve(config)# show sntp
```

SNTP Configuration

```
Time Sync Mode: Sntp
SNTP Mode : Unicast
Poll Interval (sec) [720] : 720
```

Priority	SNTP Server Address	Protocol Version
1	2001:db8::215:60ff:fe79:8980	7
2	10.255.5.24	3
3	fe80::123%vlan10	3

In this example, the **Poll Interval** and the **Protocol Version** appear at their default settings.

Both IPv4 and IPv6 addresses are displayed.

Note: Protocol Version appears only when there is an IP address configured for an SNTP server.

Figure 9-8. Example of Configuring SNTP for Unicast Operation

If the SNTP server you specify uses SNTP version 4 or later, use the **sntp server** command to specify the correct version number. For example, suppose you learned that SNTP version 4 was in use on the server you specified above (IP address 10.28.227.141). You would use the following commands to delete the server IP address and then re-enter it with the correct version number for that server:

```
ProCurve(config)# no sntp server 10.28.227.141
ProCurve(config)# sntp server 10.28.227.141 4
ProCurve(config)# show sntp
```

SNTP Configuration

```
Time Sync Mode: Sntp
SNTP Mode : Broadcast
Poll Interval (sec) [720] : 600
```

IP Address	Protocol Version
10.28.227.141	4

Deletes unicast SNTP server entry.

Re-enters the unicast server with a non-default protocol version.

show sntp displays the result.

Figure 9-9. Example of Specifying the SNTP Protocol Version Number

Changing the SNTP Poll Interval.

Syntax: `sntp poll-interval < 30..720 >`

Specifies how long the switch waits between time polling intervals. The default is 720 seconds and the range is 30 to 720 seconds. (This parameter is separate from the poll interval parameter used for Timep operation.)

For example, to change the poll interval to 300 seconds:

```
ProCurve(config)# sntp poll-interval 300
```

Changing the Priority. You can choose the order in which configured servers are polled for getting the time by setting the server priority.

Syntax: `sntp server priority <1 - 3> <ip-address>`

Specifies the order in which the configured servers are polled for getting the time. Value is between 1 and 3.

Note: Both IPv4 and IPv6 addresses can be entered. For more information about IPv6 addresses, see the “IPv6 Configuration Guide” for your switch.

For example, to set one server to priority 1 and another to priority 2:

```
ProCurve(config)# sntp server priority 1 10.28.22.141
ProCurve(config)# sntp server priority 2
                    2001:db8::215:60ff:fe79:8980
```

Disabling Time Synchronization Without Changing the SNTP Configuration. The recommended method for disabling time synchronization is to use the **timesync** command.

Syntax: `no timesync`

Halts time synchronization without changing your SNTP configuration.

For example, suppose SNTP is running as the switch’s time synchronization protocol, with **Broadcast** as the SNTP mode and the factory-default polling interval. You would halt time synchronization with this command:

```
ProCurve(config)# no timesync
```

If you then viewed the SNTP configuration, you would see the following:

```

ProCurve(config)# show sntp
SNTP Configuration
  Time Sync Mode: Disabled
  SNTP Mode : Broadcast
  Poll Interval (sec) [720] : 720

```

Figure 9-10. Example of SNTP with Time Synchronization Disabled

Disabling the SNTP Mode. If you want to prevent SNTP from being used even if selected by **timesync** (or the Menu interface's **Time Sync Method** parameter), configure the SNTP mode as disabled.

Syntax: no sntp

*Disables SNTP by changing the SNTP mode configuration to **Disabled**.*

For example, if the switch is running SNTP in Unicast mode with an SNTP server at 10.28.227.141 and a server version of 3 (the default), **no sntp** changes the SNTP configuration as shown below, and disables time synchronization on the switch.

```

ProCurve(config)# no sntp
ProCurve(config)# show sntp
SNTP Configuration
  Time Sync Mode: Sntp
  SNTP Mode : disabled
  Poll Interval (sec) [720] : 720
  IP Address          Protocol Version
  -----
  10.28.227.141      3

```

Even though the **Time Sync Mode** is set to **Sntp**, time synchronization is disabled because **no sntp** has disabled the **SNTP Mode** parameter.

Figure 9-11. Example of Disabling Time Synchronization by Disabling the SNTP Mode

SNTP Client Authentication

Enabling SNTP authentication allows network devices such as HP ProCurve switches to validate the SNTP messages received from an NTP or SNTP server before updating the network time. NTP or SNTP servers and clients must be configured with the same set of authentication keys so that the servers can authenticate the messages they send and clients (HP ProCurve switches) can validate the received messages before updating the time.

This feature provides support for SNTP client authentication on HP ProCurve switches, which addresses security considerations when deploying SNTP in a network.

Requirements

The following must be configured to enable SNTP client authentication on the switch.

SNTP Client Authentication Support

- Timesync mode must be SNTP. Use the **timesync sntp** command. (SNTP is disabled by default.)
- SNTP must be in unicast or broadcast mode. See “Configuring Unicast and Broadcast Mode” on page 9-21.
- The MD5 authentication mode must be selected.
- An SNTP authentication key-identifier (**key-id**) must be configured on the switch and a value (**key-value**) must be provided for the authentication key. A maximum of 8 sets of **key-id** and **key-value** can be configured on the switch.
- Among the keys that have been configured, one key or a set of keys must be configured as trusted. Only trusted keys will be used for SNTP authentication.
- If the SNTP server requires authentication, one of the trusted keys has to be associated with the SNTP server.
- SNTP client authentication must be enabled on the ProCurve switch. If client authentication is disabled, packets are processed without authentication. All of the above steps are necessary to enable authentication on the client.

SNTP Server Authentication Support

Note

SNTP server is not supported on ProCurve products.

The following must be performed on the SNTP server:

- The same authentication key-identifier, trusted key, authentication mode and key-value that were configured on the SNTP client must also be configured on the SNTP server.
- SNTP server authentication must be enabled on the server.

If any of the parameters on the server are changed, the parameters have to be changed on all the SNTP clients in the network as well. The authentication check will fail on the clients otherwise, and the SNTP packets will be dropped.

Configuring the Key-Identifier, Authentication Mode, and Key-Value

This command configures the **key-id**, **authentication-mode**, and **key-value**, which are required for authentication. It is executed in the global configuration context.

Syntax: sntp authentication key-id <key-id> authentication-mode <md5>
key-value <key-string> [trusted] no sntp authentication key-id <key-id>

Configures a key-id, authentication-mode (MD5 only), and key-value, which are required for authentication.

*The **no** version of the command deletes the authentication key.*

Default: No default keys are configured on the switch.

key-id: *A numeric key identifier in the range of 1-4,294,967,295 (2^{32}) that identifies the unique key value. It is sent in the SNTP packet.*

key-value <key-string>: *The secret key that is used to generate the message digest. Up to 32 characters are allowed for <key-string>.*

```
ProCurve(config)# sntp authentication key-id 55 authentication-mode md5  
key-value secretkey1
```

Figure 9-12. Example of Setting Parameters for SNTP Authentication

Configuring a Trusted Key

Trusted keys are used in SNTP authentication. In unicast mode, a **trusted** key must be associated with a specific NTP/SNTP server. That key is used for authenticating the SNTP packet.

In unicast mode, a specific server is configured on the switch so that the SNTP client communicates with the specified server to get the date and time.

In broadcast mode, the SNTP client switch checks the size of the received packet to determine if it is authenticated. If the broadcast packet is authenticated, the key-id value is checked to see if the same key-id value is configured on the SNTP client switch. If the switch is configured with the same key-id value and the key-id value is configured as “trusted”, the authentication succeeds. Only trusted key-id value information is used for SNTP authentication. See “Configuring Unicast and Broadcast Mode” on page 9-21 for information about configuring these modes.

If the packet contains key-id value information that is not configured on the SNTP client switch or the received packet contains no authentication information, it is discarded. The SNTP client switch expects packets to be authenticated if SNTP authentication is enabled.

When authentication succeeds, the time in the packet is used to update the time on the switch.

Enter the following command to configure a **key-id** as **trusted**.

Syntax: sntp authentication key-id <key-id> trusted
no sntp authentication key-id <key-id> trusted

*Trusted keys are used during the authentication process. The switch can be configured with up to eight sets of key-id/key-value pairs. One specific set must be selected for authentication; this is done by configuring the set as **trusted**.*

*The **key-id** itself must already be configured on the switch. To enable authentication, at least one **key-id** must be configured as **trusted**.*

*The **no** version of the command indicates the key is unreliable (not trusted).*

Default: No key is trusted by default.

Associating a Key with an SNTP Server

After a key is configured, it must be associated with a specific server.

Syntax: [no] sntp server priority <1-3> <ip-address | ipv6-address> <version-num> [key-id <1-4,294,967,295>]

*Configures a **key-id** to be associated with a specific server. The key itself must already be configured on the switch.*

*The **no** version of the command disassociates the key from the server. This does not remove the authentication key.*

Default: No key is associated with any server by default.

priority: *Specifies the order in which the configured servers are polled for getting the time. Value is between 1 and 3.*

<version-num>: *Specifies the SNTP software version to use, and is assigned on a per-server basis. The version setting is backwards-compatible. For example, using version 3 means that the switch accepts versions 1 through 3.*

Default: 3; range: 1 - 7.

key-id: *Optional command. The key identifier (range 1-4,294,967,295) sent in the SNTP packet. This **key-id** will be associated with the SNTP server specified in the command.*

```
ProCurve(config)# sntp server priority 1 10.10.19.5 2 key-id 55
```

Figure 9-13. Example of Associating a Key-Id with a Specific Server

Enabling SNTP Client Authentication

The **sntp authentication** command enables SNTP client authentication on the switch. If SNTP authentication is not enabled, SNTP packets are not authenticated.

Syntax: [no] sntp authentication

Enables the SNTP client authentication.

*The **no** version of the command disables authentication.*

Default: SNTP client authentication is disabled by default.

Configuring Unicast and Broadcast Mode

To enable authentication, either unicast or broadcast mode must be configured. When authentication is enabled, changing the mode from unicast to broadcast or vice versa is not allowed. You must disable authentication and then change the mode.

To set the SNTP mode or change from one mode to the other, enter the appropriate command.

Syntax: sntp unicast
sntp broadcast

Enables SNTP for either broadcast or unicast mode.

*Default: SNTP mode is disabled by default. SNTP does not operate even if specified by the CLI **timesync** command or by the menu interface **Time Sync Method** parameter.*

Unicast: *Directs the switch to poll a specific server periodically for SNTP time synchronization. The default value between each polling request is 720 seconds but can be configured. At least one manually configured server IP address is required.*

Note: *At least one **key-id** must be configured as **trusted** and it must be associated with one of the SNTP servers. To edit or remove the associated **key-id** information or SNTP server information, SNTP authentication must be disabled.*

Broadcast: *Directs the switch to acquire its time synchronization from data broadcast by any SNTP server to the network broadcast address. The switch uses the first server detected and ignores any others. However, if the Poll Interval (configurable up to 720 seconds) expires three times without the switch detecting a time update from the original server, the switch accepts a broadcast time update from the next server it detects.*

Displaying SNTP Configuration Information

The **show sntp** command displays SNTP configuration information, including any SNTP authentication keys that have been configured on the switch.

Time Protocols

SNTP: Viewing, Selecting, and Configuring

```
ProCurve(config)# show sntp

SNTP Configuration

SNTP Authentication : Enabled
Time Sync Mode: Sntp
SNTP Mode : Unicast
Poll Interval (sec) [720] : 720

Priority  SNTP Server Address                Protocol Version  KeyId
-----
1         10.10.10.2                             3                 55
2         fe80::200:24ff:fec8:4ca8                 3                 55
```

Figure 9-14. Example of SNTP Configuration Information

To display all the SNTP authentication keys that have been configured on the switch, enter the **show sntp authentication** command.

```
ProCurve(config)# show sntp authentication

SNTP Authentication Information

SNTP Authentication : Enabled

Key-ID   Auth Mode   Trusted
-----
55       MD5         Yes
10       MD5         No
```

Figure 9-15. Example of show sntp authentication Command Output

To display the statistical information for each SNTP server, enter the **sntp statistics** command. The number of SNTP packets that have failed authentication is displayed for each SNTP server address.

```
ProCurve(config)# show sntp statistics
SNTP Statistics

Received Packets : 0
Sent Packets     : 3
Dropped Packets  : 0

SNTP Server Address          Auth Failed Pkts
-----
10.10.10.1                   0
fe80::200:24ff:fec8:4ca8    0
```

Figure 9-16. Example of SNTP Authentication Statistical Information

Saving Configuration Files and the Include-Credentials Command

You can use the **include-credentials** command to store security information in the running-config file. This allows you to upload the file to a TFTP server and then later download the file to the ProCurve switches on which you want to use the same settings. For more information about the **include-credentials** command, see “Configuring Username and Password Security” in the *Access Security Guide* for your switch.

The authentication key values are shown in the output of the **show running-config** and **show config** commands only if the **include-credentials** command was executed.

When SNTP authentication is configured and **include-credentials** has not been executed, the SNTP authentication configuration is not saved.

Time Protocols

SNTP: Viewing, Selecting, and Configuring

```
ProCurve(config)# show config

Startup configuration:
.
.
.
timesync sntp
sntp broadcast
sntp 50
sntp authentication
sntp server priority 1 10.10.10.2 3 key-id 55
sntp server priority 2 fe80::200:24ff:fec8:4ca8 4 key-id 55
.
.
.
```

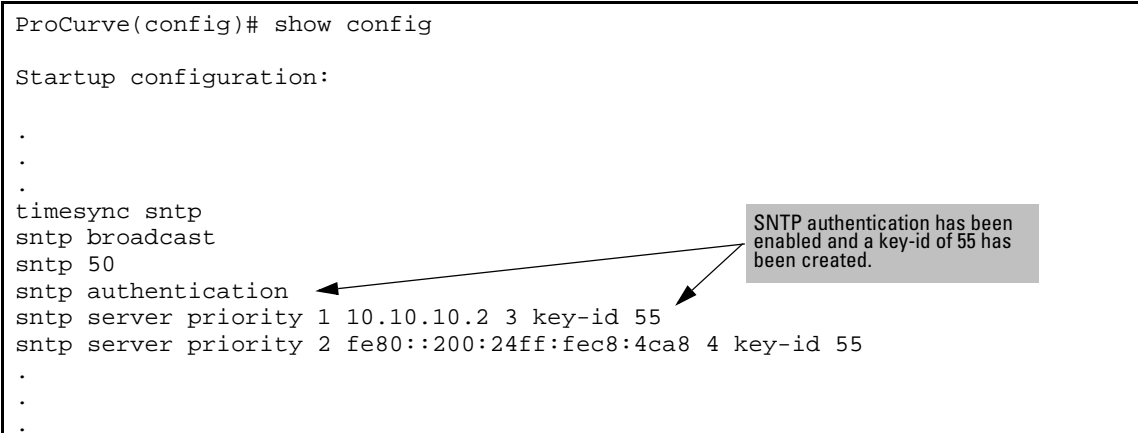


Figure 9-17. Example of Configuration File with SNTP Authentication Information

In figure 9-17, the **include-credentials** command has not been executed and is not present in the configuration file. The configuration file is subsequently saved to a TFTP server for later use. The SNTP authentication information is not saved and is not present in the retrieved configuration file, as shown in figure 9-18.

```
ProCurve(config)#copy tftp startup-config 10.2.3.44 config1
.
.
.
Switch reboots...

Startup configuration
.
.
.
timesync sntp
sntp broadcast
sntp 50 sntp server priority 1 10.10.10.2 3
sntp server priority 2 fe80::200:24ff:fec8:4ca8 4
.
.
.
```

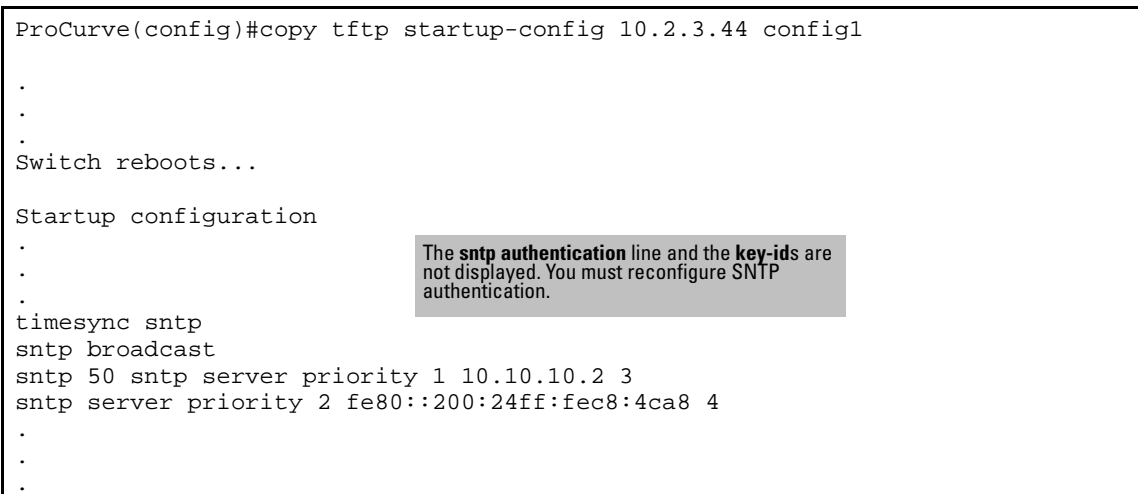


Figure 9-18. Example of a Retrieved Configuration File When Include Credentials is not Configured

If **include-credentials** is configured, the SNTP authentication configuration is saved in the configuration file. When the **show config** command is entered, all of the information that has been configured for SNTP authentication displays, including the key-values.

```
ProCurve(config)# show config

Startup configuration:

.
.
.
include-credentials
timesync sntp
sntp broadcast
sntp 50
sntp authentication
sntp authentication key-id 55 authentication-mode md5 key-value "secretkey1"
trusted
sntp authentication key-id 2 authentication-mode md5 key-value "secretkey2"
sntp server priority 1 10.10.10.2 3 key-id 55
sntp server priority 2 fe80::200:24ff:fec8:4ca8 4 key-id 55
sntp server priority 3 10.10.4.60 3
.
.
.
```

Figure 9-19. Example of Saved SNTP Authentication Information when include-credentials is Configured

TimeP: Viewing, Selecting, and Configuring

TimeP Feature	Default	Menu	CLI	Web
view the Timep time synchronization configuration	n/a	page 9-27	page 9-29	—
select Timep as the time synchronization method	TIMEP	page 9-16	pages 9-31 ff.	—
disable time synchronization	timep	page 9-27	page 9-33	—
enable the Timep mode	Disabled			—
DHCP	—	page 9-27	page 9-31	—
manual	—	page 9-28	page 9-32	—
none/disabled	—	page 9-27	page 9-34	—
change the SNTP poll interval	720 minutes	page 9-28	page 9-33	—

Table 9-2. Timep Parameters

SNTP Parameter	Operation
Time Sync Method	Used to select either TIMEP (the default), SNTP, or None as the time synchronization method.
Timep Mode	
Disabled	The Default. Timep does not operate, even if specified by the Menu interface Time Sync Method parameter or the CLI timesync command.
DHCP	When Timep is selected as the time synchronization method, the switch attempts to acquire a Timep server IP address via DHCP. If the switch receives a server address, it polls the server for updates according to the Timep poll interval. If the switch does not receive a Timep server IP address, it cannot perform time synchronization updates.
Manual	When Timep is selected as the time synchronization method, the switch attempts to poll the specified server for updates according to the Timep poll interval. If the switch fails to receive updates from the server, time synchronization updates do not occur.
Server Address	Used only when the TimeP Mode is set to Manual . Specifies the IP address of the TimeP server that the switch accesses for time synchronization updates. You can configure one server.

Menu: Viewing and Configuring TimeP

To View, Enable, and Modify the TimeP Protocol:

1. From the Main Menu, select:

2. Switch Configuration...

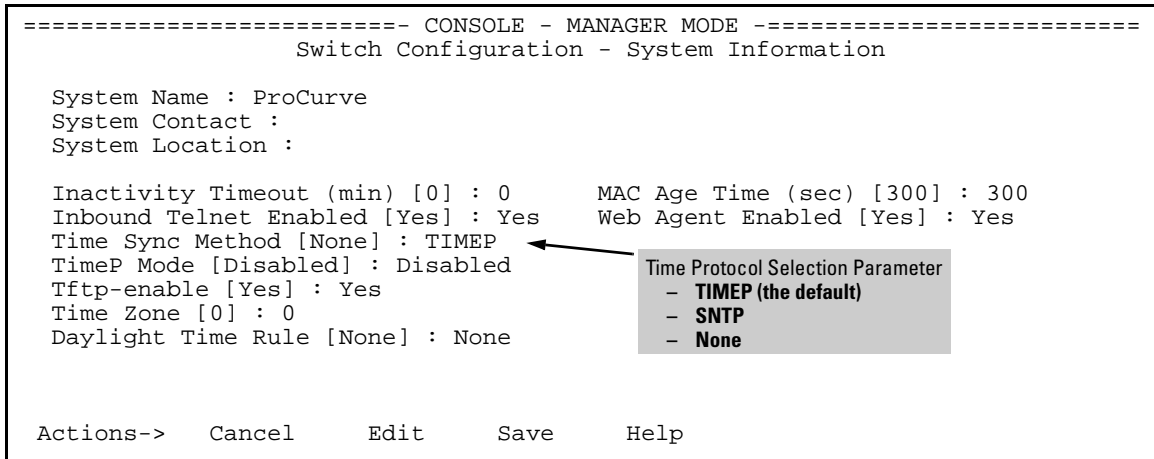
1. System Information

```
===== CONSOLE - MANAGER MODE =====
Switch Configuration - System Information

System Name : ProCurve
System Contact :
System Location :

Inactivity Timeout (min) [0] : 0      MAC Age Time (sec) [300] : 300
Inbound Telnet Enabled [Yes] : Yes    Web Agent Enabled [Yes] : Yes
Time Sync Method [None] : TIMEP
TimeP Mode [Disabled] : Disabled
Tftp-enable [Yes] : Yes
Time Zone [0] : 0
Daylight Time Rule [None] : None

Actions->  Cancel      Edit      Save      Help
```



The screenshot shows a terminal window titled "CONSOLE - MANAGER MODE" with the subtitle "Switch Configuration - System Information". It displays various system parameters and their current values. A callout box on the right, titled "Time Protocol Selection Parameter", lists three options: "TIMEP (the default)", "SNTP", and "None". An arrow points from this callout box to the "TimeP Mode [Disabled] : Disabled" line in the terminal output.

Figure 9-20. The System Information Screen (Default Values)

Press **[E]** (for **E**dit). The cursor moves to the **System Name** field.

2. Use **[↓]** to move the cursor to the **Time Sync Method** field.

3. If **TIMEP** is not already selected, use the Space bar to select **TIMEP**, then press **[↓]** once to display and move to the **TimeP Mode** field.

4. Do one of the following:

- Use the Space bar to select the **DHCP** mode, then press **[↓]** to move the cursor to the **Poll Interval** field, and go to step 6.

Time Protocols

TimeP: Viewing, Selecting, and Configuring

```
Time Sync Method [None] : TIMEP
TimeP Mode [Disabled] : DHCP
Poll Interval (min) [720] : 720
Time Zone [0] : 0
Daylight Time Rule [None] : None
```

- Use the Space bar to select the **Manual** mode.
 - i. Press `→` to move the cursor to the **Server Address** field.
 - ii. Enter the IP address of the TimeP server you want the switch to use for time synchronization.

Note: This step replaces any previously configured TimeP server IP address.
 - iii. Press `→` to move the cursor to the **Poll Interval** field, then go to step 6.
- 5. In the **Poll Interval** field, enter the time in minutes that you want for a TimeP Poll Interval.

Press **[Enter]** to return to the Actions line, then **[S]** (for **Save**) to enter the new time protocol configuration in both the startup-config and running-config files.

CLI: Viewing and Configuring TimeP

CLI Commands Described in this Section

Command	Page
show timep	9-29
[no] timesync	9-30 ff., 9-33
ip timep	
dhcp	9-31
manual	9-32
server <ip-addr>	9-32
interval	9-33
no ip timep	9-34

Viewing the Current TimeP Configuration

Using different **show** commands, you can display either the full TimeP configuration or a combined listing of all TimeP, SNTP, and VLAN IP addresses configured on the switch.

Syntax: show timep

*This command lists both the time synchronization method (TimeP, SNTP, or None) and the TimeP configuration, even if SNTP is not the selected time protocol. (If the TimeP Mode is set to **Disabled** or **DHCP**, then the **Server** field does not appear.)*

For example, if you configure the switch with TimeP as the time synchronization method, then enable TimeP in DHCP mode with the default poll interval, **show timep** lists the following:

```
ProCurve(config)# show timep
Timep Configuration
Time Sync Mode: Timep
TimeP Mode [Disabled] : DHCP      Server Address : 10.10.28.100
Poll Interval (min) [720] : 720
```

Figure 9-21. Example of TimeP Configuration When TimeP Is the Selected Time Synchronization Method

If SNTP is the selected time synchronization method, **show timep** still lists the TimeP configuration even though it is not currently in use:

```
ProCurve(config)# show timep
Timep Configuration
Time Sync Mode: Sntp
[TimeP Mode [Disabled] : Manual      Server Address : 10.10.28.100]
[Poll Interval (min) [720] : 720]
```

Even though, in this example, SNTP is the current time synchronization method, the switch maintains the TimeP configuration.

Figure 9-22. Example of TimeP Configuration When TimeP Is Not the Selected Time Synchronization Method

Syntax: show management

This command can help you to easily examine and compare the IP addressing on the switch. It lists the IP addresses for all time servers configured on the switch, plus the IP addresses and default gateway for all VLANs configured on the switch.

Time Protocols

TimeP: Viewing, Selecting, and Configuring

```
ProCurve(config)# show management

Status and Counters - Management Address Information

Time Server Address : 10.10.28.100

Priority  SNTP Server Address                Protocol Version
-----  -
1         10.10..28.101                        3
2         10.255.5.24                          3
3         fe80::123%vlan10                    3

Default Gateway      : 10.0.9.80

VLAN Name      MAC Address      | IP Address
-----  -
DEFAULT_VLAN  001279-88a100    | 10.30.248.184
VLAN10        001279-88a100    | 10.0.10.17
```

Figure 9-23. Example of Display Showing IP Addressing for All Configured Time Servers and VLANs

Configuring (Enabling or Disabling) the TimeP Mode

Enabling the TimeP mode means to configure it for either broadcast or unicast mode. Remember that to run TimeP as the switch's time synchronization protocol, you must also select TimeP as the time synchronization method by using the CLI `timesync` command (or the Menu interface **Time Sync Method** parameter).

Syntax: `timesync timep`
Selects TimeP as the time protocol.

Syntax: `ip timep < dhcp | manual >`
Enables the selected TimeP mode.

Syntax: `no ip timep`
Disables the TimeP mode.

Syntax: `no timesync`
Disables the time protocol.

Enabling TimeP in DHCP Mode. Because the switch provides a TimeP polling interval (default: 720 minutes), you need only these two commands for a minimal TimeP DHCP configuration:

Syntax: `timesync timep`

Selects TimeP as the time synchronization method.

Syntax: `ip timep dhcp`

Configures DHCP as the TimeP mode.

For example, suppose:

- Time synchronization is configured for SNTP.
- You want to:
 1. View the current time synchronization.
 2. Select TimeP as the time synchronization mode.
 3. Enable TimeP for DHCP mode.
 4. View the TimeP configuration.

The commands and output would appear as follows:

```
ProCurve(config)# show timep ❶ show timep displays the TimeP configuration and also shows
Timep Configuration that SNTP is the currently active time synchronization mode.
Time Sync Mode: Sntp
TimeP Mode : Disabled

ProCurve(config)# timesync timep ❷

ProCurve(config)# ip timep dhcp ❸

ProCurve(config)# show timep ❹ show timep again displays the TimeP configuration and shows that TimeP is
Timep Configuration now the currently active time synchronization mode.
Time Sync Mode: Timep
TimeP Mode : DHCP Poll Interval (min) : 720
```

Figure 9-24. Example of Enabling TimeP Operation in DHCP Mode

Enabling TimeP in Manual Mode. Like DHCP mode, configuring TimeP for **Manual** mode enables TimeP. However, for manual operation, you must also specify the IP address of the TimeP server. (The switch allows only one TimeP server.) To enable the TimeP protocol:

Syntax: timesync timep

Selects TimeP.

Syntax: ip timep manual < ip-addr > [oobm]

Activates TimeP in Manual mode with a specified TimeP server.

*For switches that have a separate out-of-band management port, **oobm** specifies that SNTP traffic goes through that port. (By default, SNTP traffic goes through the data ports.)*

Syntax: no ip timep

Disables TimeP.

Note

To change from one TimeP server to another, you must (1) use the **no ip timep** command to disable TimeP mode, and then reconfigure TimeP in Manual mode with the new server IP address.

For example, to select TimeP and configure it for manual operation using a TimeP server address of 10.28.227.141 and the default poll interval (720 minutes, assuming the TimeP poll interval is already set to the default):

```
ProCurve(config)# timesync timep
```

Selects TimeP.

```
ProCurve(config)# ip timep manual 10.28.227.141
```

Activates TimeP in Manual mode.

```
ProCurve(config)# timesync timep
ProCurve(config)# ip timep manual 10.28.227.141

ProCurve(config)# Show timep
Timep Configuration
  Time Sync Mode: Timep
  TimeP Mode : Manual           Server Address : 10.28.227.141
  Poll Interval (min) : 720
```

Figure 9-25. Example of Configuring Timep for Manual Operation

Changing the TimeP Poll Interval. This command lets you specify how long the switch waits between time polling intervals. The default is 720 minutes and the range is 1 to 9999 minutes. (This parameter is separate from the poll interval parameter used for SNTP operation.)

Syntax: ip timep < dhcp | manual > interval < 1 - 9999 >

For example, to change the poll interval to 60 minutes:

```
ProCurve(config)# ip timep interval 60
```

Disabling Time Synchronization Without Changing the TimeP Configuration. The recommended method for disabling time synchronization is to use the **timesync** command. This halts time synchronization without changing your TimeP configuration.

Syntax: no timesync

*Disables time synchronization by changing the **Time Sync Mode** configuration to **Disabled**.*

For example, suppose TimeP is running as the switch's time synchronization protocol, with **DHCP** as the TimeP mode, and the factory-default polling interval. You would halt time synchronization with this command:

```
HPswitch(config)# no timesync
```

If you then viewed the TimeP configuration, you would see the following:

```
ProCurve(config)# show timep
Timep Configuration
Time Sync Mode: Disabled
TimeP Mode : DHCP    Poll Interval (min) : 720
```

Figure 9-26. Example of TimeP with Time Synchronization Disabled

Disabling the TimeP Mode. Disabling the TimeP mode means to configure it as disabled. (Disabling TimeP prevents the switch from using it as the time synchronization protocol, even if it is the selected **Time Sync Method** option.)

Syntax: no ip timep

*Disables TimeP by changing the TimeP mode configuration to **Disabled**.*

For example, if the switch is running TimeP in DHCP mode, **no ip timep** changes the TimeP configuration as shown below, and disables time synchronization.

```
ProCurve(config)# no ip timep

ProCurve(config)# show timep
Timep Configuration
Time Sync Mode: Timep
TimeP Mode : Disabled
```

Even though the Time Sync Mode is set to Timep, time synchronization is disabled because no ip timep has disabled the TimeP Mode parameter.

Figure 9-27. Example of Disabling Time Synchronization by Disabling the TimeP Mode Parameter

SNTP Unicast Time Polling with Multiple SNTP Servers

When running SNTP unicast time polling as the time synchronization method, the switch requests a time update from the server you configured with either the Server Address parameter in the menu interface, or the primary server in a list of up to three SNTP servers configured using the CLI. If the switch does not receive a response from the primary server after three consecutive polling intervals, the switch tries the next server (if any) in the list. If the switch tries all servers in the list without success, it sends an error message to the Event Log and reschedules to try the address list again after the configured **Poll Interval** time has expired.

Displaying All SNTP Server Addresses Configured on the Switch

The System Information screen in the menu interface displays only one SNTP server address, even if the switch is configured for two or three servers. The CLI **show management** command displays all configured SNTP servers on the switch.

```
ProCurve(config)# show management

Status and Counters - Management Address Information

Time Server Address : fe80::215:60ff:fe7a:adc0%vlan10

Priority SNTP Server Address                                Protocol Version
-----
1          2001:db8::215:60ff:fe79:8980                    7
2          10.255.5.24                                       3
3          fe80::123%vlan10                                 3

Default Gateway      : 10.0.9.80

VLAN Name      MAC Address      | IP Address
-----
DEFAULT_VLAN  001279-88a100    | Disabled
VLAN10        001279-88a100    | 10.0.10.17
```

Figure 9-28. Example of How To List All SNTP Servers Configured on the Switch

Adding and Deleting SNTP Server Addresses

Adding Addresses. As mentioned earlier, you can configure one SNTP server address using either the Menu interface or the CLI. To configure a second and third address, you must use the CLI. To configure the remaining two addresses, you would do the following:

```
ProCurve(config)# sntp server 2001:db8::215:60ff:fe79:8980
ProCurve(config)# sntp server 10.255.5.24
```

Figure 9-29. Example of Creating Additional SNTP Server Addresses with the CLI

Note

If there are already three SNTP server addresses configured on the switch, and you want to use the CLI to replace one of the existing addresses with a new one, you must delete the unwanted address before you configure the new one.

Deleting Addresses. To delete an address, you must use the CLI. If there are multiple addresses and you delete one of them, the switch re-orders the address priority.

Syntax: `no sntp server < ip-addr >`

For example, to delete the primary address in the above example (and automatically convert the secondary address to primary):

```
ProCurve(config)# no sntp server 10.28.227.141
```

Menu: Operation with Multiple SNTP Server Addresses Configured

When you use the Menu interface to configure an SNTP server IP address, the new address writes over the current primary address, if one is configured.

SNTP Messages in the Event Log

If an SNTP time change of more than three seconds occurs, the switch's event log records the change. SNTP time changes of less than three seconds do not appear in the Event Log.

Port Status and Configuration

Contents

Overview	10-3
Viewing Port Status and Configuring Port Parameters	10-3
Menu: Port Configuration	10-6
CLI: Viewing Port Status and Configuring Port Parameters	10-8
Viewing Port Status and Configuration	10-8
Customizing the Show Interfaces Command	10-10
Error Messages	10-12
Note on Using Pattern Matching with the “Show Interfaces Custom” Command	10-13
Viewing Port Utilization Statistics	10-13
Viewing Transceiver Status	10-14
Enabling or Disabling Ports and Configuring Port Mode	10-15
Enabling or Disabling the USB Port	10-17
Behavior of Autorun When USB Port is Disabled	10-18
Enabling or Disabling Flow Control	10-18
Configuring a Broadcast Limit on the Switch	10-20
Configuring ProCurve Auto-MDIX	10-21
Web: Viewing Port Status and Configuring Port Parameters	10-24
Using Friendly (Optional) Port Names	10-25
Configuring and Operating Rules for Friendly Port Names	10-25
Configuring Friendly Port Names	10-26
Displaying Friendly Port Names with Other Port Data	10-27
Configuring Transceivers and Modules That Haven’t Been Inserted	10-31
Transceivers	10-31
Modules	10-31
Clearing the Module Configuration	10-31

Operating Notes	10-32
Uni-Directional Link Detection (UDLD)	10-33
Configuring UDLD	10-34
Enabling UDLD	10-35
Changing the Keepalive Interval	10-36
Changing the Keepalive Retries	10-36
Configuring UDLD for Tagged Ports	10-36
Viewing UDLD Information	10-37
Configuration Warnings and Event Log Messages	10-39

Overview

This chapter describes how to view the current port configuration and how to configure ports to non-default settings, including

- Enable/Disable
- Mode (speed and duplex)
- Flow Control
- Broadcast Limit
- Friendly Port Names
- Uni-directional Link Detection (UDLD)

Viewing Port Status and Configuring Port Parameters

Port Status and Configuration Features

Feature	Default	Menu	CLI	Web
viewing port status	n/a	page 10-6	page 10-8	page 10-24
viewing transceiver status	n/a	n/a	page 10-14	page 10-24
configuring ports	Refer to Table 10-1 on pages 10-4 thru 10-5	page 10-7	page 10-15	page 10-24
configuring ProCurve auto-mdix			page 9-11	

Note On Connecting Transceivers to Fixed-Configuration Devices

If the switch either fails to show a link between an installed transceiver and another device, or demonstrates errors or other unexpected behavior on the link, check the port configuration on both devices for a speed and/or duplex (mode) mismatch.

- To check the mode setting for a port on the switch, use either the Port Status screen in the menu interface (page 10-6) or **show interfaces brief** in the CLI (page 10-8).

To display information about the transceivers installed on a switch, enter the **show tech receivers** command in the CLI (page 10-14).

Port Status and Configuration

Viewing Port Status and Configuring Port Parameters

Table 10-1. Status and Parameters for Each Port Type

Status or Parameter	Description
Enabled	Yes (default): The port is ready for a network connection. No : The port will not operate, even if properly connected in a network. Use this setting, for example, if the port needs to be shut down for diagnostic purposes or while you are making topology changes.
Status (read-only)	Up : The port senses a link beat. Down : The port is not enabled, has no cables connected, or is experiencing a network error. For troubleshooting information, refer to the <i>Installation and Getting Started Guide</i> you received with the switch. Refer also to Appendix C, “Troubleshooting” (in this manual).
Mode	The port’s speed and duplex (data transfer operation) setting.

10/100/1000Base-T Ports:

- **Auto-MDIX** (default): Senses speed and negotiates with the port at the other end of the link for port operation (MDI-X or MDI).
To see what the switch negotiates for the Auto setting, use the CLI **show interfaces brief** command or the “3. Port Status” option under “1. Status and Counters” in the menu interface.
- **MDI**: Sets the port to connect with a PC using a crossover cable (Manual mode—applies only to copper port switches using twisted-pair copper Ethernet cables)
- **MDIX**: Sets the port to connect with a PC using a straight-through cable (Manual mode—applies only to copper port switches using twisted-pair copper Ethernet cables)
- **Auto-10**: Allows the port to negotiate between half-duplex (HDx) and full-duplex (FDx) while keeping speed at 10 Mbps. Also negotiates flow control (enabled or disabled). ProCurve recommends Auto-10 for links between 10/100 auto-sensing ports connected with Cat 3 cabling. (Cat 5 cabling is required for 100 Mbps links.)
- **10HDx**: 10 Mbps, Half-Duplex
- **10FDx**: 10 Mbps, Full-Duplex
- **Auto-100**: Uses 100 Mbps and negotiates with the port at the other end of the link for other port operation features.
- **Auto-10-100**: Allows the port to establish a link with the port at the other end at either 10 Mbps or 100 Mbps, using the highest mutual speed and duplex mode available. Only these speeds are allowed with this setting.
- **Auto-1000**: Uses 1000 Mbps and negotiates with the port at the other end of the link for other port operation features.
- **100Hdx**: Uses 100 Mbps, half-duplex.
- **100Fdx**: Uses 100 Mbps, Full-Duplex

— Continued on Next Page —

Status or Parameter	Description
<i>— Continued From Previous Page —</i>	
	<p>Gigabit Fiber-Optic Ports (Gigabit-SX, Gigabit-LX, and Gigabit-LH):</p> <ul style="list-style-type: none"> • 1000FDx: 1000 Mbps (1 Gbps), Full Duplex only • Auto (default): The port operates at 1000FDx and auto-negotiates flow control with the device connected to the port. <p>Gigabit Copper Ports:</p> <ul style="list-style-type: none"> • 1000FDx: 1000 Mbps (1 Gbps), Full Duplex only • Auto (default): The port operates at 1000FDx and auto-negotiates flow control with the device connected to the port. <hr/> <p>10-Gigabit CX4 Copper Ports:</p> <ul style="list-style-type: none"> • Auto: The port operates at 10 gigabits FDx and negotiates flow control. Lower speed settings or half-duplex are not allowed. <p>10-Gigabit SC Fiber-Optic Ports (10-GbE SR, 10-GbE LR, 10-GbE ER):</p> <ul style="list-style-type: none"> • Auto: The port operates at 10 gigabits FDx and negotiates flow control. Lower speed settings or half-duplex are not allowed. <p>Note: Conditioning patch cord cables are not supported on 10-GbE.</p> <hr/>
Auto-MDIX	<p>The switch supports Auto-MDIX on 10Mb, 100Mb, and 1 Gb T/TX (copper) ports. (Fiber ports and 10-gigabit ports do not use this feature.)</p> <ul style="list-style-type: none"> • Automdix: Configures the port for automatic detection of the cable type (straight-through or crossover). • MDI: Configures the port to connect to a switch, hub, or other MDI-X device with a straight-through cable. • MDIX: Configures the port to connect to a PC or other MDI device with a straight-through cable. <hr/>
Flow Control	<ul style="list-style-type: none"> • Disabled (default): The port does not generate flow control packets, and drops any flow control packets it receives. • Enabled: The port uses 802.3x Link Layer Flow Control, generates flow control packets, and processes received flow control packets. <p>With the port mode set to Auto (the default) and Flow Control enabled, the switch negotiates Flow Control on the indicated port. If the port mode is not set to Auto, or if Flow Control is disabled on the port, then Flow Control is not used. Note that flow control must be enabled on both ends of a link.</p> <hr/>
Broadcast Limit	<p>Specifies the percentage of the theoretical maximum network bandwidth that can be used for broadcast and multicast traffic. Any broadcast or multicast traffic exceeding that limit will be dropped. Zero (0) means the feature is disabled.</p> <p>The broadcast-limit command operates at the port context level to set the broadcast limit for a port on the switch.</p> <p>Note: This feature is not appropriate for networks that require high levels of IPX or RIP broadcast traffic.</p> <hr/>

Menu: Port Configuration

From the menu interface, you can view and change the port configuration.

Using the Menu To View Port Configuration. The menu interface displays the configuration for ports and (if configured) any trunk groups.

From the Main Menu, select:

1. Status and Counters
4. Port Status

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

Port	Type	Intrusion Alert	Enabled	Status	Mode	MDI Mode	Flow Ctrl
B1	10/100TX	No	Yes	Down	10FDx	Auto	off
B2	10/100TX	No	Yes	Down	10FDx	Auto	off
B3	10/100TX	No	Yes	Down	10FDx	Auto	off
B4	10/100TX	No	Yes	Down	10FDx	Auto	off
B5	10/100TX	No	Yes	Down	10FDx	Auto	off
B6	10/100TX	No	Yes	Down	10FDx	Auto	off
B7-Trk2	10/100TX	No	Yes	Down	10FDx	Auto	off
B8-Trk2	10/100TX	No	Yes	Down	10FDx	Auto	off
B9	10/100TX	No	Yes	Down	10FDx	Auto	off
B10	10/100TX	No	Yes	Down	10FDx	Auto	off
B11	10/100TX	No	Yes	Down	10FDx	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.

In this example, ports A7 and A8 have previously been configured as a trunk group.



Figure 10-1. Example of a Switch Port Status Screen

Using the Menu To Configure Ports.

You can configure and view the port settings by using the menu.

Note

The menu interface uses the same screen for configuring both individual ports and port trunk groups. For information on port trunk groups, refer to Chapter 12, “Port Trunking” .

1. From the Main Menu, Select:

2. Switch Configuration...

2. Port/Trunk Settings

An example of the Menu display is shown below.

```
===== TELNET - MANAGER MODE =====
                Switch Configuration - Port/Trunk Settings

Port      Type      Enabled      Mode      Flow Ctrl  Group  Type
----      -+-----      -+-----      -+-----      -+-----      -+-----
A1      1000T      | Yes      Auto-10-100  Disable
A2      1000T      | Yes      Auto-10-100  Disable
A3      1000T      | Yes      Auto          Disable
A4      1000T      | Yes      Auto          Disable
A5      1000T      | Yes      Auto          Disable
A6      1000T      | Yes      Auto          Disable
A7      1000T      | Yes      Auto          Disable      Trk1  Trunk
A8      1000T      | Yes      Auto          Disable      Trk2  Trunk

Actions->  Cancel      Edit      Save      Help

Cancel changes and return to previous screen.
Use arrow keys to change action selection and <Enter> to execute ac-
tion.
```

Figure 10-2. Example of Port/Trunk Settings with a Trunk Group Configured

2. Press **[E]** (for **E**dit). The cursor moves to the **Enabled** field for the first port.
3. Refer to the online help provided with this screen for further information on configuration options for these features.
4. When you have finished making changes to the above parameters, press **[Enter]**, then press **[S]** (for **S**ave).

CLI: Viewing Port Status and Configuring Port Parameters

From the CLI, you can configure and view all port parameter settings and view all port status indicators.

Port Status and Configuration Commands

show interfaces brief	page 10-9
show interfaces config	page 10-9
show interfaces custom	page 10-10
show interfaces port-utilization	page 10-13
show tech transceivers	page 10-14
interface	page 10-15
disable/enable	page 10-15
speed-duplex	page 10-15
flow-control	page 10-18
broadcast-limit	page 10-20
auto-mdix	page 10-21

Viewing Port Status and Configuration

Use the following commands to display port status and configuration data.

Syntax: show interfaces [brief | config | < port-list >]

brief: Lists the current operating status for all ports on the switch.

config: Lists a subset of configuration data for all ports on the switch; that is, for each port, the display shows whether the port is enabled, the operating mode, and whether it is configured for flow control.

< port-list >: Shows a summary of network traffic handled by the specified ports.

An example of the **show interfaces brief** command is shown below.

```
ProCurve(config)# show interfaces brief
Status and Counters - Port Status
```

Port	Type	Intrusion		Status	Mode	MDI Mode	Flow Ctrl	Bcast Limit
		Alert	Enabled					
B1	100/1000T	No	Yes	Down	Auto-10-100	Auto	off	0
B2	100/1000T	No	Yes	Down	1000FDx	Auto	off	0
B3	100/1000T	No	Yes	Down	1000FDx	Auto	off	0
B4	100/1000T	No	Yes	Down	1000FDx	Auto	off	0
B5	100/1000T	No	Yes	Down	1000FDx	Auto	off	0
B6	100/1000T	No	Yes	Down	1000FDx	Auto	off	0

Figure 10-3. Example of Show Interfaces Brief Command Listing

Use the **show interfaces config** command to view the port settings, as shown below.

```
ProCurve(config)# show interfaces config

Port Settings
```

Port	Type	Enabled	Mode	Flow Ctrl	MDI
B1	100/1000T	Yes	Auto-10-100	Disable	Auto
B2	100/1000T	Yes	Auto	Disable	Auto
B3	100/1000T	Yes	Auto	Disable	Auto
B4	100/1000T	Yes	Auto	Disable	Auto
B5	100/1000T	Yes	Auto	Disable	Auto
B6	100/1000T	Yes	Auto	Disable	Auto

Figure 10-4. Example of a Show Interfaces Config Command Listing

The **display** option can be used to initiate the dynamic update of the **show interfaces** command with the output being the same as the **show interfaces** command. When using the **display** option in the CLI, the information stays on the screen and is updated every 3 seconds, as occurs with the display using the menu feature. The update is terminated with Cntl-C.

You can use the arrow keys to scroll through the screen when the output does not fit in one screen.

Syntax: show interfaces display

Initiates the dynamic update of a command. The output is the same as the equivalent “show” command. The information is updated every 3 seconds.

Note: Select “Back” to exit the display.

For example:

ProCurve# show interfaces display

Port	Total Bytes	Total Frames	Errors Rx	Drops Tx	Flow Ctrl	Bca*
1	2,164,277	20,366	0	0	off	0
2	0	0	0	0	off	0
3	0	0	0	0	off	0
4	0	0	0	0	off	0
5	0	0	0	0	off	0
6	0	0	0	0	off	0
7	0	0	0	0	off	0
8	0	0	0	0	off	0
9	0	0	0	0	off	0
10	0	0	0	0	off	0
11	0	0	0	0	off	0

Actions-> **Back** Show details Reset Help

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

Figure 10-5. Example of show interfaces display Command with Dynamically Updating Output

Customizing the Show Interfaces Command

You can create **show** commands displaying the information that you want to see in any order you want by using the **custom** option.

Syntax: show interfaces custom [port-list] column-list

Select the information that you want to display. Parameters include:

- *port name*
- *type*
- *vlan*
- *intrusion*
- *enabled*
- *status*
- *speed*
- *mdi*
- *flow*

Columns supported are:

Parameter Column	Displays	Examples
port	Port identifier	A2
type	Port type	100/100T
status	Port status	up or down
speed	Connection speed and duplex	1000FDX
mode	Configured mode	auto, auto-100, 100FDX
mdi	MDI mode	auto, MDIX
flow	Flow control	on or off
name	Friendly port name	
vlanid	The vlan id this port belongs to, or "tagged" if it belongs to more than one vlan	4 tagged
enabled	port is or is not enabled	yes or no intrusion
intrusion	Intrusion alert status	no
bcast	Broadcast limit	0

```
ProCurve(config)# show int custom 1-4 port name:4 type vlan intrusion speed
enabled mdi

Status and Counters - Custom Port Status

Port Name      Type      VLAN  Intrusion
Alert         Speed    Enabled MDI-mode
-----
1   Acco     100/1000T  1   No      1000FDx  Yes   Auto
2   Huma     100/1000T  1   No      1000FDx  Yes   Auto
3   Deve     100/1000T  1   No      1000FDx  Yes   Auto
4   Lab1     100/1000T  1   No      1000FDx  Yes   Auto
```

Figure 10-6. Example of the Custom show interfaces Command

You can specify the column width by entering a colon after the column name, then indicating the number of characters to display. In Figure 10-6 the Name column only displays the first four characters of the name. All remaining characters are truncated.

Note

Each field has an fixed minimum width to be displayed. If you specify a field width smaller than the minimum width, the information is displayed at the minimum width. For example, if the minimum width for the Name field is 4 characters and you specify Name:2, the Name field displays 4 characters.

Parameters can be entered in any order. There is a limit of 80 characters per line; if you exceed this limit an error displays.

Error Messages

Error	Error Message
Requesting too many fields (total characters exceeds 80)	Total length of selected data exceeds one line
Field name is misspelled	Invalid input: <input>
Mistake in specifying the port list	Module not present for port or invalid port: <input>
The port list is not specified	Incomplete input: custom

Note on Using Pattern Matching with the “Show Interfaces Custom” Command

If you have included a pattern matching command to search for a field in the output of the **show int custom** command and the **show int custom** command produces an error, the error message may not be visible and the output is empty. For example, if you enter a command that produces an error (vlan is misspelled) with the pattern matching **include** option:

```
ProCurve(config)# show int custom 1-3 name vlun |  
include vlan1
```

the output may be empty. It is advisable to try the **show int custom** command first to ensure there is output, and then enter the command again with the pattern matching option.

Viewing Port Utilization Statistics

Use the **show interface port-utilization** command to view a real-time rate display for all ports on the switch. The following shows a sample output from this command.

```
ProCurve(config)# show interfaces port-utilization  
Status and Counters - Port Utilization
```

Port	Mode	Rx			Tx		
		Kbits/sec	Pkts/sec	Util	Kbits/sec	Pkts/sec	Util
B1	1000FDx	0	0	0	0	0	0
B2	1000FDx	0	0	0	0	0	0
B3	1000FDx	0	0	0	0	0	0
B4	1000FDx	0	0	0	0	0	0
B5	1000FDx	0	0	0	0	0	0
B6	1000FDx	0	0	0	0	0	0
B7	100FDx	624	86	00.62	496	0	00.49

Figure 10-7. Example of a Show Interface Port-Utilization Command Listing

Operating Notes:

- For each port on the switch, the command provides a real-time display of the rate at which data is received (Rx) and transmitted (Tx) in terms of kilobits per second (KBits/s), number of packets per second (Pkts/s), and utilization (Util) expressed as a percentage of the total bandwidth available.
- The **show interfaces** <port-list> command can be used to display the current link status and the port rate average over a 5 minute period. Port rates are shown in bits per second (bps) for ports up to 1 Gigabit; for 10 Gigabit ports, port rates are shown in kilobits per second (Kbps).

Viewing Transceiver Status

The **show tech transceivers** command allows you to:

- Remotely identify transceiver type and revision number without having to physically remove an installed transceiver from its slot.
- Display real-time status information about all installed transceivers, including non-operational transceivers.

Figure 10-8 shows sample output from the **show tech transceivers** command.

```
ProCurve# show tech transceivers

Transceiver Technical Information:
Port # | Type      | Prod # | Serial #      | Part #
-----+-----+-----+-----+-----
21    | 1000SX   | J4858B | CN605MP23K   |
22    | 1000LX   | J4859C | H117E7X      | 2157-2345
23    | ??       | ??     | non operational |
25    | 10GbE-CX4 | J8440A | US509RU079   |
26    | 10GbE-CX4 | J8440A | US540RU002   |
27    | 10GbE-LR | J8437B | PPA02-2904:0017 | 2157-2345
28    | 10GbE-SR | J8436B | 01591602     | 2158-1000
29    | 10GbE-ER | J8438A | PPA03-2905:0001 |

The following transceivers may not function correctly:
Port #      Message
-----
Port 23     Self test failure.
```

Figure 10-8. Example of Show Tech Transceivers Command

Operating Notes:

- The following information is displayed for each installed transceiver:
 - Port number on which transceiver is installed.
 - Type of transceiver.
 - Product number—Includes revision letter, such as A, B, or C. If no revision letter follows a product number, this means that no revision is available for the transceiver.
 - Part number—Allows you to determine the manufacturer for a specified transceiver and revision number.
- For a non-ProCurve installed transceiver (see line 23 Figure 10-8), no transceiver type, product number, or part information is displayed. In the Serial Number field, **non-operational** is displayed instead of a serial number.
- The following error messages may be displayed for a non-operational transceiver:
 - `Unsupported Transceiver. (SelfTest Err#060)`
`Check: www.hp.com/rnd/device_help/2_inform` for more info.
 - `This switch only supports revision B and above transceivers. Check: www.hp.com/rnd/device_help/2_inform` for more info.
 - `Self test failure.`
 - `Transceiver type not supported in this port.`
 - `Transceiver type not supported in this software version.`
 - `Not a ProCurve Transceiver. Please go to: www.hp.com/rnd/device_help/2_inform` for more info.

Enabling or Disabling Ports and Configuring Port Mode

You can configure one or more of the following port parameters. Refer to table 10-1 on pages 10-4 through 10-5.

Syntax: `[no] interface <port-list>`
`[<disable | enable>]`

*Disables or enables the port for network traffic. Does not use the **no** form of the command. (Default: **enable**.)*

```
speed-duplex < auto-10 | 10-full | 10-half | 100-full | 100-half | auto | auto-100 | 1000-full >]
```

*Specifies the port's data transfer speed and mode. Does not use the **no** form of the command. (Default: **auto**.)*

Note that in the above syntax you can substitute an “**int**” for “**interface**”; that is: **int < port-list >**.

The 10/100 auto-negotiation feature allows a port to establish a link with a port at the other end at either 10 Mbps or 100 Mbps, using the highest mutual speed and duplex mode available. Only these speeds are allowed with this setting.

For example, to configure port C5 for auto-10-100, enter this command:

```
ProCurve(config)# int c5 speed-duplex auto-10-100
```

To configure ports C1 through C3 and port C6 for 100Mbps full-duplex, you would enter these commands:

```
ProCurve(config)# int c1-c3,c6 speed-duplex 100-full
```

Similarly, to configure a single port with the above command settings, you could either enter the same command with only the one port identified, or go to the *context level* for that port and then enter the command. For example, to enter the context level for port C6 and then configure that port for 100FDx:

```
ProCurve(config)# int e c6
ProCurve(eth-C6)# speed-duplex 100-full
```

If port C8 was disabled, and you wanted to enable it and configure it for 100FDx with flow-control active, you could do so with either of the following command sets.

```
ProCurve(config)# int c8 enable
ProCurve(config)# int c8 speed-duplex 100-full
ProCurve(config)# int c8 flow-control

ProCurve(config)# int c8
ProCurve(eth-C8)# enable
ProCurve(eth-C8)# speed-duplex 100-full
```

These commands enable and configure port C8 from the config level:

These commands select the port C8 context level and then apply the subsequent configuration commands to port C8:

Figure 10-9. Examples of Two Methods for Changing a Port Configuration

Refer to “Enabling or Disabling Flow Control” on page 10-18 for more on flow control.

Enabling or Disabling the USB Port

This feature allows configuration of the USB port with either the CLI or SNMP.

To enable/disable the USB port with the CLI:

Syntax: usb-port
no usb-port

*Enables the USB port. The **no** form of the command disables the USB port and any access to the device.*

To display the status of the USB port:

Syntax: show usb-port

Displays the status of the USB port. It can be enabled, disabled, or not present.

```
ProCurve(config)# show usb-port
USB port status: enabled
USB port power status: power on      (USB device detected in port)
USB port reseal status: USB reseal not required
```

Figure 10-10. Example of show usb-port Command Output on version K.13.59 and later

```
ProCurve(config)# show usb-port
USB port status: enabled
USB port power status: power on      (USB device detected in port)
```

Figure 10-11. Example of show usb-port Command Output on version K.14.XX

One of the following messages indicates the presence or absence of the USB device:

- Not able to sense device in USB port
- USB device detected in port
- no USB device detected in port

The reseal status messages can be one of the following (K.13.XX only):

- undetermined USB reseal requirement
- USB reseal not required
- USB device reseal required for USB autorun

The autorun feature only works when a USB device is inserted and the USB port is enabled.

Behavior of Autorun When USB Port is Disabled

Software Versions K.13.XX Operation.

When using software version K.13.58, if the USB port is disabled (no `usb-port` command), the USB autorun function does not work in the USB port until the USB port is enabled, the config file is saved, and the switch is rebooted. The 5 volt power to the USB port remains on even after the USB port has been disabled.

For software versions after K.13.58, the 5 volt power applied to the USB port is synchronized with the enabling of the USB port, that is, when the USB port is enabled, the 5 volts are supplied; when the USB port is disabled, the 5 volts are not supplied. For previous software versions the power was supplied continuously. The autorun function does not require a switch reboot, but the USB device must be inserted at least once after the port is enabled so that the switch recognizes that the device is present. If the USB device is inserted and then the USB port is enabled, the switch does not recognize that a USB device is present.

Software Version K.14.XX Operation.

For software versions K.14.XX, the USB port can be disabled and enabled without affecting the autorun feature. When the USB port is enabled, the autorun feature activates if a USB device is already inserted in the USB port.

Power is synchronized with the enabling and disabling of USB ports as described above for K.13.59 and later software.

Enabling or Disabling Flow Control

Note

You must enable flow control on both ports in a given link. Otherwise, flow control does not operate on the link, and appears as **Off** in the **show interfaces brief** port listing, even if flow control is configured as enabled on the port in the switch. (Refer to Figure 10-3 on page 10-9.) Also, the port (speed-duplex) mode must be set to **Auto** (the default).

To disable flow control on some ports, while leaving it enabled on other ports, just disable it on the individual ports you want to exclude.

Syntax: [no]interface < port-list > flow-control

Enables or disables flow control packets on the port. The “no” form of the command disables flow control on the individual ports. (Default: Disabled.)

For example, suppose that:

1. You want to enable flow control on ports A1-A6.
2. Later, you decide to disable flow control on ports A5 and A6.
3. As a final step, you want to disable flow control on all ports.

Assuming that flow control is currently disabled on the switch, you would use these commands:

```
ProCurve(config)# int a1-a6 flow-control
ProCurve(config)# show interfaces brief
Status and Counters - Port Status
```

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

← Enables per-port flow control for ports A1 - A6.

Figure 10-12. Example of Configuring Flow Control for a Series of Ports

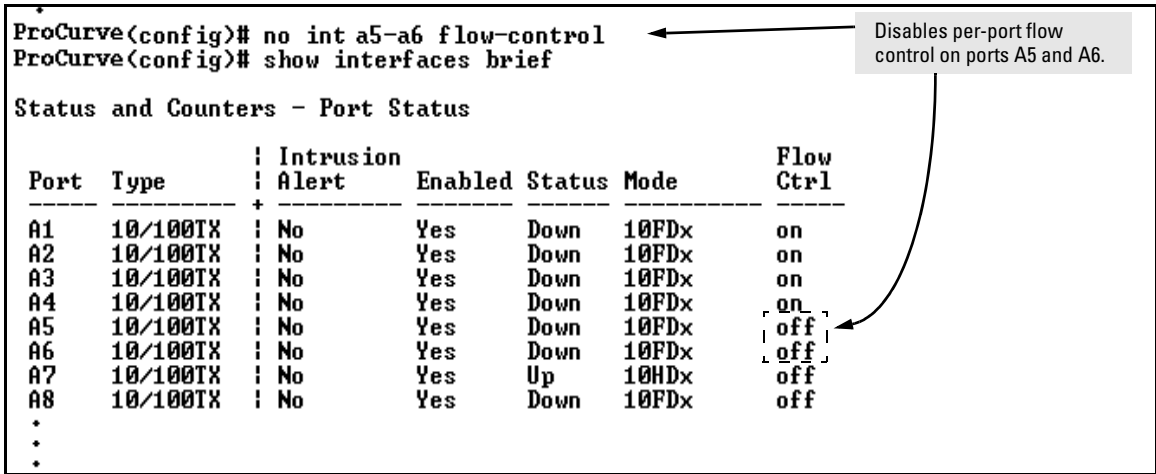


Figure 10-13. Example Continued from Figure 10-12

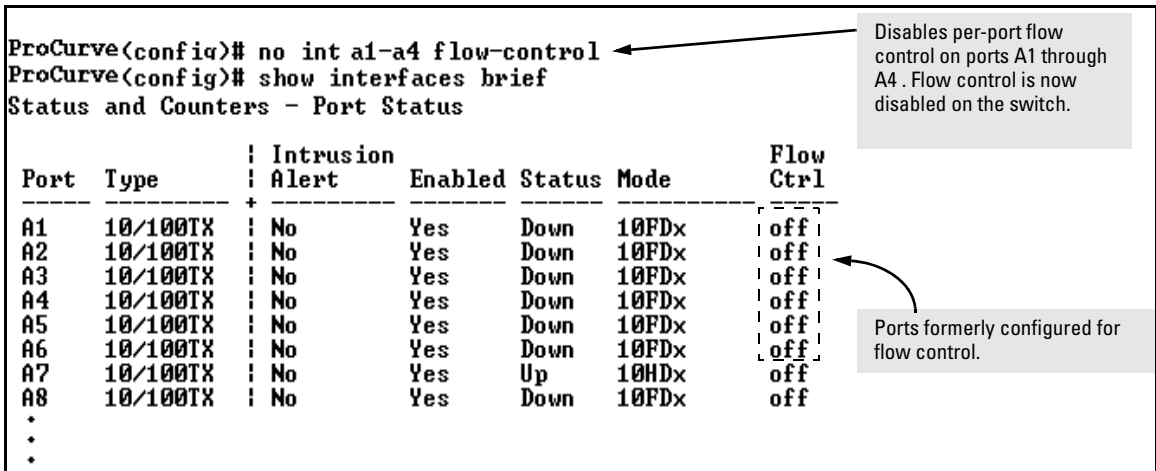


Figure 10-14. Example Continued from Figure 10-13

Configuring a Broadcast Limit on the Switch

Broadcast-Limit on switches covered in this guide is configured on a per-port basis. You must be at the port context level for this command to work, for example:

```

ProCurve(config)#int B1
ProCurve(int B1)# broadcast-limit 1

```

Syntax: broadcast-limit <0-99>

*Enables or disables broadcast limiting for outbound broadcasts on a selected port on the switch. The value selected is the percentage of traffic allowed, for example, **broadcast-limit 5** allows 5% of the maximum amount of traffic for that port. A value of zero disables broadcast limiting for that port.*

Note: You must switch to port context level before issuing the **broadcast-limit** command.

Note: This feature is not appropriate for networks requiring high levels of IPX or RIP broadcast traffic.

Syntax: show config

Displays the startup-config file. The broadcast limit setting appears here if enabled and saved to the startup-config file.

Syntax: show running-config

Displays the running-config file. The broadcast limit setting appears here if enabled. If the setting is not also saved to the startup-config file, rebooting the switch returns broadcast limit to the setting currently in the startup-config file.

For example, the following command enables broadcast limiting of 1 percent of the traffic rate on the selected port on the switch:

```
ProCurve(int B1)# broadcast-limit 1
```

For a one Gbps port this results in a broadcast traffic rate of ten Mbps.

Configuring ProCurve Auto-MDIX

Copper ports on the switch can automatically detect the type of cable configuration (MDI or MDI-X) on a connected device and adjust to operate appropriately.

This means you can use a “straight-through” twisted-pair cable or a “cross-over” twisted-pair cable for any of the connections—the port makes the necessary adjustments to accommodate either one for correct operation. The following port types on your switch support the IEEE 802.3ab standard, which includes the “Auto MDI/MDI-X” feature:

- 10/100-TX xl module ports
- 100/1000-T xl module ports
- 10/100/1000-T xl module ports

Using the above ports:

- If you connect a copper port using a straight-through cable on a switch to a port on another switch or hub that uses MDI-X ports, the switch port automatically operates as an MDI port.
- If you connect a copper port using a straight-through cable on a switch to a port on an end node, such as a server or PC, that uses MDI ports, the switch port automatically operates as an MDI-X port.

ProCurve Auto-MDIX was developed for auto-negotiating devices, and was shared with the IEEE for the development of the IEEE 802.3ab standard. ProCurve Auto-MDIX and the IEEE 802.3ab Auto MDI/MID-X feature are completely compatible. Additionally, ProCurve Auto-MDIX supports operation in forced speed and duplex modes.

If you want more information on this subject please refer to the *IEEE 802.3ab Standard Reference*.

For more information on MDI-X, refer to the appendix titled “Switch Ports and Network Cables” in the *Installation and Getting Started Guide* for your switch.

Manual Override. If you require control over the MDI/MDI-X feature you can set the switch to either of two non-default modes:

- Manual MDI
- Manual MDI-X

Table 10-2 shows the cabling requirements for the MDI/MDI-X settings.

Table 10-2. Cable Types for Auto and Manual MDI/MDI-X Settings

Setting	MDI/MDI-X Device Type	
	PC or Other MDI Device Type	Switch, Hub, or Other MDI-X Device
Manual MDI	Crossover Cable	Straight-Through Cable
Manual MDI-X	Straight-Through Cable	Crossover Cable
Auto-MDI-X (The Default)	Either Crossover or Straight-Through Cable	

The Auto-MDIX features apply only to copper port switches using twisted-pair copper Ethernet cables.

Syntax: interface < port-list > mdix-mode < auto-mdix | mdi | mdix >

auto-mdix is the automatic, default setting. This configures the port for automatic detection of the cable (either straight-through or crossover).

mdi is the manual mode setting that configures the port for connecting to either a PC or other MDI device with a crossover cable, or to a switch, hub, or other MDI-X device with a straight-through cable.

mdix is the manual mode setting that configures the port for connecting to either a switch, hub, or other MDI-X device with a crossover cable, or to a PC or other MDI device with a straight-through cable.

Syntax: show interfaces config

Lists the current per-port Auto/MDI/MDI-X configuration.

Syntax: show interfaces brief

*Where a port is linked to another device, this command lists the MDI mode the port is currently using. In the case of ports configured for **Auto (auto-mdix)**, the MDI mode appears as either **MDI** or **MDIX**, depending upon which option the port has negotiated with the device on the other end of the link. In the case of ports configured for **MDI** or **MDIX**, the mode listed in this display matches the configured setting. If the link to another device was up, but has gone down, this command shows the last operating MDI mode the port was using. If a port on a given switch has not detected a link to another device since the last reboot, this command lists the MDI mode to which the port is currently configured.*

For example, **show interfaces config** displays the following data when port A1 is configured for **auto-mdix**, port A2 is configured for **mdi**, and port A3 is configured for **mdix**.

```
ProCurve(config)# show interfaces config
```

Port Settings						Per-Port MDI Configuration
Port	Type	Enabled	Mode	Flow Ctrl	MDI	
A1	10/100TX	Yes	Auto	Disable	Auto	↖
A2	10/100TX	Yes	Auto	Disable	MDI	
A3	10/100TX	Yes	Auto	Disable	MDIX	
A4	10/100TX	Yes	Auto	Disable	Auto	
A5	10/100TX	Yes	Auto	Disable	Auto	
.	
.	

Figure 10-15. Example of Displaying the Current MDI Configuration

```
ProCurve(config)# show interfaces brief
```

Status and Counters - Port Status								Per-Port MDI Operating Mode
Port	Type	Intrusion Alert	Enabled	Status	Mode	MDI Mode	Flow Ctrl	
A1	10/100TX	No	Yes	Up	100FDx	MDIX	off	↖
A2	10/100TX	No	Yes	Up	100FDx	MDI	off	
A3	10/100TX	No	Yes	Up	100FDx	MDIX	off	
A4	10/100TX	No	Yes	Down	10FDx	Auto	off	
A5	10/100TX	No	Yes	Down	10FDx	Auto	off	
.	
.	

Figure 10-16. Example of Displaying the Current MDI Operating Mode

Web: Viewing Port Status and Configuring Port Parameters

In the web browser interface:

1. Click on the **Configuration** tab.
2. Click on **[Port Configuration]**.
3. Select the ports you want to modify and click on **[Modify Selected Ports]**.
4. After you make the desired changes, click on **[Apply Settings]**.

Note that the web browser interface displays an existing port trunk group. However, to configure a port trunk group, you must use the CLI or the menu interface. For more on this topic, refer to Chapter 12, “Port Trunking” .

Using Friendly (Optional) Port Names

Feature	Default	Menu	CLI	Web
Configure Friendly Port Names	Standard Port Numbering	n/a	page 26	n/a
Display Friendly Port Names	n/a	n/a	page 27	n/a

This feature enables you to assign alphanumeric port names of your choosing to augment automatically assigned numeric port names. This means you can configure meaningful port names to make it easier to identify the source of information listed by some **Show** commands. (Note that this feature *augments* port numbering, but *does not replace* it.)

Configuring and Operating Rules for Friendly Port Names

- At either the global or context configuration level you can assign a unique name to a port. You can also assign the same name to multiple ports.
- The friendly port names you configure appear in the output of the **show name [port-list]**, **show config**, and **show interface <port-number>** commands. They do not appear in the output of other show commands or in Menu interface screens. (Refer to “Displaying Friendly Port Names with Other Port Data” on page 10-27.)
- Friendly port names are not a substitute for port numbers in CLI commands or Menu displays.
- Trunking ports together does not affect friendly naming for the individual ports. (If you want the same name for all ports in a trunk, you must individually assign the name to each port.)
- A friendly port name can have up to 64 contiguous alphanumeric characters.
- Blank spaces within friendly port names are not allowed, and if used, cause an **invalid input** error. (The switch interprets a blank space as a name terminator.)
- In a port listing, **not assigned** indicates that the port does not have a name assignment other than its fixed port number.

- To retain friendly port names across reboots, you must save the current running-configuration to the startup-config file after entering the friendly port names. (In the CLI, use the **write memory** command.)

Configuring Friendly Port Names

Syntax: interface < port-list > name < port-name-string >
Assigns a port name to port-list.

Syntax: no interface < port-list > name
Deletes the port name from port-list.

Configuring a Single Port Name. Suppose that you have connected port A3 on the switch to Bill Smith's workstation, and want to assign Bill's name and workstation IP address (10.25.101.73) as a port name for port A3:

```
ProCurve(config)# int A3 name Bill_Smith@10.25.101.73
ProCurve(config)# write mem
ProCurve(config)# show name A3

Port Names
Port : A3
Type : 10/100TX
Name : Bill_Smith@10.25.101.73
```

Figure 10-17. Example of Configuring a Friendly Port Name

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

```
ProCurve(config)# int A5-A8 name Draft-Server:Trunk
ProCurve(config)# write mem
ProCurve(config)# show name 5-8

Port Names

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

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

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

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

Figure 10-18. Example of Configuring One Friendly Port Name on Multiple Ports

Displaying Friendly Port Names with Other Port Data

You can display friendly port name data in the following combinations:

- **show name:** Displays a listing of port numbers with their corresponding friendly port names and also quickly shows you which ports do not have friendly name assignments. (**show name** data comes from the running-config file.)
- **show interface <port-number>:** Displays the friendly port name, if any, along with the traffic statistics for that port. (The friendly port name data comes from the running-config file.)
- **show config:** Includes friendly port names in the per-port data of the resulting configuration listing. (**show config** data comes from the startup-config file.)

To List All Ports or Selected Ports with Their Friendly Port Names.

This command lists names assigned to a specific port.

Syntax: show name [port-list]

*Lists the friendly port name with its corresponding port number and port type. The **show name** command without a port list shows this data for all ports on the switch.*

For example:

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

Figure 10-19. Example of Friendly Port Name Data for All Ports on the Switch

```
ProCurve(config)# show name A2,A3,A5
Port Names
Port : A2
Type : 10/100TX
Name : not assigned
Port : A3
Type : 10/100TX
Name : Bill_Smith@10.25.101.73
Port : A5
Type : 10/100TX
Name : Draft-Server:Trunk
```

Figure 10-20. Example of Friendly Port Name Data for Specific Ports on the Switch

Including Friendly Port Names in Per-Port Statistics Listings. A friendly port name configured to a port is automatically included when you display the port's statistics output.

Syntax: show interface < port-number >

Includes the friendly port name with the port's traffic statistics listing.

For example, if you configure port A1 with the name "O'Connor_10.25.101.43", the show interface output for this port appears similar to the following:

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

Name      : O'Connor@10.25.101.43 ← Friendly Port Name

Link Status      : Up

Bytes Rx         : 894,568           Bytes Tx         : 2470
Unicast Rx      : 1179              Unicast Tx       : 13
Bcast/Mcast Rx  : 5280              Bcast/Mcast Tx   : 13

FCS Rx          : 36                Drops Tx         : 0
Alignment Rx    : 2                 Collisions Tx    : 0
Runts Rx        : 0                 Late Colln Tx   : 0
Giants Rx       : 0                 Excessive Colln : 0
Total Rx Errors : 38                Deferred Tx      : 0
```

Figure 10-21. Example of a Friendly Port Name in a Per-Port Statistics Listing

For a given port, if a friendly port name does not exist in the running-config file, the Name line in the above command output appears as:

```
Name      : not assigned
```

To Search the Configuration for Ports with Friendly Port Names.

This option tells you which friendly port names have been saved to the startup-config file. (**show config** does not include ports that have only default settings in the startup-config file.)

Syntax: show config

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

Port Status and Configuration

Using Friendly (Optional) Port Names

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

```
ProCurve(config)# int A1 name Print_Server@10.25.101.43
ProCurve(config)# write mem
ProCurve(config)# int A2 name Herbert's_PC
ProCurve(config)# show config

Startup configuration:
; J4850A Configuration Editor; Created on release #E.08.30
hostname "HPswitch"
time daylight-time-rule None
no cdp run
interface A1
  name "Print_Server@10.25.101.43"
exit
snmp-server community "public" Unrestricted
vlan 1
  name "DEFAULT_VLAN"
  untagged 1-24
  ip address dhcp-bootp
  exit
no aaa port-access authenticator active
```

This command sequence saves the friendly port name for port A1 in the startup-config file. The name entered for port A2 is not saved because it was executed after **write memory**.

Listing includes friendly port name for port A1 only.

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

Figure 10-22. Example Listing of the Startup-Config File with a Friendly Port Name Configured (and Saved)

Configuring Transceivers and Modules That Haven't Been Inserted

Transceivers

Previously, a port had to be valid and verified for the switch to allow it to be configured. Transceivers are removable ports and considered invalid when not present in the switch, so they cannot be configured unless they are already in the switch. For switches covered in this guide, the verification for allowable port configurations performed by the CLI is removed and configuration of transceivers is allowed even if they are not yet inserted in the switch.

Modules

You can create or edit configuration files (as text files) that can be uploaded to the switch without the modules having been installed yet. Additionally, you can pre-configure the modules with the CLI “**module**” command.

Syntax: `module <module-num> type <module-type>`

Allows you to configure the type of the module.

The same **module** command used in an uploaded configuration file is used to define a module that is being pre-configured. The validation performed when issued through the CLI is still performed just as if the command was executed on the switch, in other words, as if the module were actually present in the switch.

Note

You cannot use this method to change the configuration of a module that has already been configured. The slot must be empty and the configuration file must not have a configuration associated with it.

Clearing the Module Configuration

Because of the hot-swap capabilities of the modules, when a module is removed from the chassis, the module configuration remains in the configuration file. This feature allows you to remove the module configuration information from the configuration file.

Syntax: [no] module <slot>

Allows removal of the module configuration in the configuration file after the module has been removed. Enter an integer between 1 and 12 for <slot>.

For example:

```
ProCurve(config)# no module 3
```

Note

This does not change how hot-swap works.

Operating Notes

The following restrictions apply:

- The slot being cleared must be empty
- There was no module present in the slot since the last boot
- If there was a module present after the switch was booted, the switch will have to be rebooted before any module (new or same) can be used in the slot.
- This does not clear the configuration of a module still in use by the switch.

Uni-Directional Link Detection (UDLD)

Uni-directional Link Detection (UDLD) monitors a link between two ProCurve switches and blocks the ports on both ends of the link if the link fails at any point between the two devices. This feature is particularly useful for detecting failures in fiber links and trunks. Figure 10-23 shows an example.

Scenario 1 (No UDLD): Without UDLD, the switch ports remain enabled despite the link failure. Traffic continues to be load-balanced to the ports connected to the failed link.

Scenario 2 (UDLD-enabled): When UDLD is enabled, the feature blocks the ports connected to the failed link.

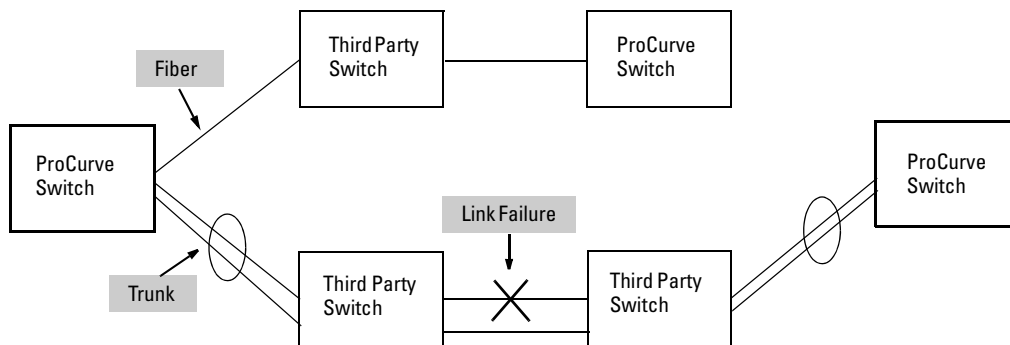


Figure 10-23. UDLD Example

In this example, each ProCurve switch load balances traffic across two ports in a trunk group. Without the UDLD feature, a link failure on a link that is not directly attached to one of the ProCurve switches remains undetected. As a result, each switch continues to send traffic on the ports connected to the failed link. When UDLD is enabled on the trunk ports on each ProCurve switch, the switches detect the failed link, block the ports connected to the failed link, and use the remaining ports in the trunk group to forward the traffic.

Similarly, UDLD is effective for monitoring fiber optic links that use two uni-directional fibers to transmit and receive packets. Without UDLD, if a fiber breaks in one direction, a fiber port may assume the link is still good (because the other direction is operating normally) and continue to send traffic on the

connected ports. UDLD-enabled ports; however, will prevent traffic from being sent across a bad link by blocking the ports in the event that either the individual transmitter or receiver for that connection fails.

Ports enabled for UDLD exchange health-check packets once every five seconds (the link-keepalive interval). If a port does not receive a health-check packet from the port at the other end of the link within the keepalive interval, the port waits for four more intervals. If the port still does not receive a health-check packet after waiting for five intervals, the port concludes that the link has failed and blocks the UDLD-enabled port.

When a port is blocked by UDLD, the event is recorded in the switch log or via an SNMP trap (if configured); and other port blocking protocols, like spanning tree or meshing, will not use the bad link to load balance packets. The port will remain blocked until the link is unplugged, disabled, or fixed. The port can also be unblocked by disabling UDLD on the port.

Configuring UDLD

When configuring UDLD, keep the following considerations in mind:

- UDLD is configured on a per-port basis and must be enabled at both ends of the link. See the note below for a list of ProCurve switches that support UDLD.
- To configure UDLD on a trunk group, you must configure the feature on each port of the group individually. Configuring UDLD on a trunk group's primary port enables the feature on that port only.
- Dynamic trunking is not supported. If you want to configure a trunk group that contains ports on which UDLD is enabled, you must remove the UDLD configuration from the ports. After you create the trunk group, you can re-add the UDLD configuration.

Note

UDLD interoperates with the following ProCurve switches: 2600, 2800, 3400, 3500, 4200, 5300, 5400, 6200, 6400, 6600, 8212, and 9300. Consult the release notes and current manuals for required software versions.

The following commands allow you to configure UDLD via the CLI.

Syntax: [no] interface <port-list> link-keepalive

Enables UDLD on a port or range of ports.

*To disable the feature, enter the **no** form of the command.*

Default: UDLD disabled

Syntax: link-keepalive interval <interval>

Determines the time interval to send UDLD control packets. The <interval> parameter specifies how often the ports send a UDLD packet. You can specify from 10 – 100, in 100 ms increments, where 10 is 1 second, 11 is 1.1 seconds, and so on.

Default: 50 (5 seconds)

Syntax: link-keepalive retries <num>

Determines the maximum number of retries to send UDLD control packets. The <num> parameter specifies the maximum number of times the port will try the health check. You can specify a value from 3 – 10.

Default: 5

Syntax: [no] interface <port-list> link-keepalive vlan <vid>

Assigns a VLAN ID to a UDLD-enabled port for sending of tagged UDLD control packets. Under default settings, untagged UDLD packets can still be transmitted and received on tagged only ports—however, a warning message will be logged.

*The **no** form of the command disables UDLD on the specified port(s).*

Default: UDLD packets are untagged; tagged only ports will transmit and receive untagged UDLD control packets

Enabling UDLD

UDLD is enabled on a per port basis. For example, to enable UDLD on port a1, enter:

```
ProCurve(config)#interface a1 link-keepalive
```

To enable the feature on a trunk group, enter the appropriate port range. For example:

```
ProCurve(config)#interface a1-a4 link-keepalive
```

Note

When at least one port is UDLD-enabled, the switch will forward out UDLD packets that arrive on non-UDLD-configured ports out of all other non-UDLD-configured ports in the same vlan. That is, UDLD control packets will “pass through” a port that is not configured for UDLD. However, UDLD packets will be dropped on any blocked ports that are not configured for UDLD.

Changing the Keepalive Interval

By default, ports enabled for UDLD send a link health-check packet once every 5 seconds. You can change the interval to a value from 10 – 100 deciseconds, where 10 is 1 second, 11 is 1.1 seconds, and so on. For example, to change the packet interval to seven seconds, enter the following command at the global configuration level:

```
ProCurve(config)# link-keepalive interval 70
```

Changing the Keepalive Retries

By default, a port waits five seconds to receive a health-check reply packet from the port at the other end of the link. If the port does not receive a reply, the port tries four more times by sending up to four more health-check packets. If the port still does not receive a reply after the maximum number of retries, the port goes down.

You can change the maximum number of keepalive attempts to a value from 3 – 10. For example, to change the maximum number of attempts to 4, enter the following command at the global configuration level:

```
ProCurve(config)# link-keepalive retries 4
```

Configuring UDLD for Tagged Ports

The default implementation of UDLD sends the UDLD control packets untagged, even across tagged ports. If an untagged UDLD packet is received by a non-ProCurve switch, that switch may reject the packet. To avoid such an occurrence, you can configure ports to send out UDLD control packets that are tagged with a specified VLAN.

To enable ports to receive and send UDLD control packets tagged with a specific VLAN ID, enter a command such as the following at the interface configuration level:

```
ProCurve(config)#interface 1 link-keepalive vlan 22
```

Notes

- You must configure the same VLANs that will be used for UDLD on all devices across the network; otherwise, the UDLD link cannot be maintained.
- If a VLAN ID is not specified, then UDLD control packets are sent out of the port as untagged packets.

- To re-assign a VLAN ID, re-enter the command with the new VLAN ID number. The new command will overwrite the previous command setting.
- When configuring UDLD for tagged ports, you may receive a warning message if there are any inconsistencies with the port's VLAN configuration (see page 39 for potential problems).

Viewing UDLD Information

The following show commands allow you to display UDLD configuration and status via the CLI.

Syntax: show link-keepalive

Displays all the ports that are enabled for link-keepalive.

Syntax: show link-keepalive statistics

Displays detailed statistics for the UDLD-enabled ports on the switch.

Syntax: clear link-keepalive statistics

Clears UDLD statistics. This command clears the packets sent, packets received, and transitions counters in the show link-keepalive statistics display.

To display summary information on all UDLD-enabled ports, enter the **show link-keepalive** command. For example:

```
ProCurve(config)# show link-keepalive

Total link-keepalive enabled ports: 4
Keepalive Retries: 3           Keepalive Interval: 1 sec

Port Enabled Physical  Keepalive  Adjacent  UDLD
      Status Status      Status    Switch    VLAN
-----
1  Yes  up      up      00d9d-f9b700  200
2  Yes  up      up      01560-7b1600
3  Yes  down    off-line
4  Yes  up      failure
5  No   down    off-line
```

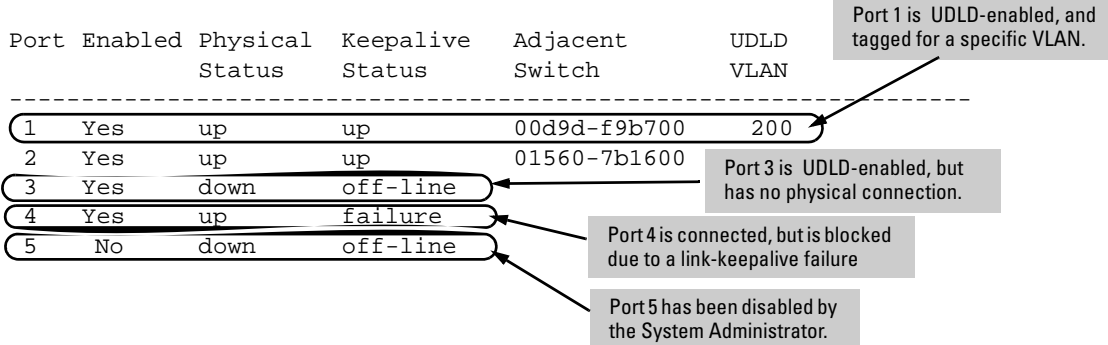


Figure 10-24. Example of Show Link-Keepalive Command

Port Status and Configuration

Uni-Directional Link Detection (UDLD)

To display detailed UDLD information for specific ports, enter the **show link-keepalive statistics** command. For example:

```
ProCurve(config)# show link-keepalive statistics
```

Port:	1		
Current State:	up	Neighbor MAC Addr:	0000a1-b1c1d1
Uddl Packets Sent:	1000	Neighbor Port:	5
Uddl Packets Received:	1000	State Transitions:	2
Port Blocking:	no	Link-vlan:	1
Port:	2		
Current State:	up	Neighbor MAC Addr:	000102-030405
Uddl Packets Sent:	500	Neighbor Port:	6
Uddl Packets Received:	450	State Transitions:	3
Port Blocking:	no	Link-vlan:	200
Port:	3		
Current State:	off line	Neighbor MAC Addr:	n/a
Uddl Packets Sent:	0	Neighbor Port:	n/a
Uddl Packets Received:	0	State Transitions:	0
Port Blocking:	no	Link-vlan:	1
Port:	4		
Current State:	failure	Neighbor MAC Addr:	n/a
Uddl Packets Sent:	128	Neighbor Port:	n/a
Uddl Packets Received:	50	State Transitions:	8
Port Blocking:	yes	Link-vlan:	1

Ports 1 and 2 are UDLD-enabled and show the number of health check packets sent and received on each port.

Port 4 is shown as blocked due to a link-keepalive failure

Figure 10-25. Example of Show Link-Keepalive Statistics Command

To clear UDLD statistics, enter the following command:

```
ProCurve# clear link-keepalive statistics
```

This command clears the packets sent, packets received, and transitions counters in the **show link keepalive statistics** display (see Figure 10-25 for an example).

Configuration Warnings and Event Log Messages

Warning Messages. The following table shows the warning messages that may be issued and their possible causes, when UDLD is configured for tagged ports.

Table 10-3. Warning Messages caused by configuring UDLD for Tagged Ports

CLI Command Example	Warning Message	Possible Problem
link-keepalive 6	Possible configuration problem detected on port 6. UDLD VLAN configuration does not match the port's VLAN configuration.	You have attempted to enable UDLD on a port that is a tagged only port, but did not specify a configuration for tagged UDLD control packets. In this example, the switch will send and receive the UDLD control packets untagged despite issuing this warning.
link-keepalive 7 vlan 4	Possible configuration problem detected on port 7. UDLD VLAN configuration does not match the port's VLAN configuration.	You have attempted to configure tagged UDLD packets on a port that does not belong to the specified VLAN. In this example, if port 7 belongs to VLAN 1 and 22, but the user tries to configure UDLD on port 7 to send tagged packets in VLAN 4, the configuration will be accepted. The UDLD control packets will be sent tagged in VLAN 4, which may result in the port being blocked by UDLD if the user does not configure VLAN 4 on this port.
no vlan 22 tagged 20	Possible configuration problem detected on port 18. UDLD VLAN configuration does not match the port's VLAN configuration.	You have attempted to remove a VLAN on port that is configured for tagged UDLD packets on that VLAN. In this example, if port 18, 19, and 20 are transmitting and receiving tagged UDLD packets for Vlan 22, but the user tries to remove Vlan 22 on port 20, the configuration will be accepted. In this case, the UDLD packets will still be sent on Vlan 20, which may result in the port being blocked by UDLD if the users do not change the UDLD configuration on this port.

Note: If you are configuring the switch via SNMP with the same problematic VLAN configuration choices, the above warning messages will also be logged in the switch's event log.

Event Log Messages. The following table shows the event log messages that may be generated once UDLD has been enabled on a port.

Table 10-4. UDLD Event Log Messages

Message	Event
I 01/01/06 04:25:05 ports: port 4 is deactivated due to link failure.	A UDLD-enabled port has been blocked due to part of the link having failed.
I 01/01/06 06:00:43 ports: port 4 is up, link status is good.	A failed link has been repaired and the UDLD-enabled port is no longer blocked.

Port Status and Configuration
Uni-Directional Link Detection (UDLD)

Power Over Ethernet (PoE/PoE+) Operation

Contents

Introduction to PoE	11-2
PoE Terminology	11-2
PoE Operation	11-4
Configuration Options	11-4
PD Support	11-5
Power Priority Operation	11-6
When Is Power Allocation Prioritized?	11-6
How Is Power Allocation Prioritized?	11-6
Configuring PoE Operation	11-7
Disabling or Re-Enabling PoE Port Operation	11-7
Enabling Support for Pre-Standard Devices	11-7
Configuring the PoE Port Priority Level	11-8
PoE Priority With Two or More Modules	11-9
Controlling PoE Allocation	11-11
Manually Configuring PoE Power Levels	11-12
Configuring PoE Redundancy (Chassis Switches Only)	11-13
Changing the Threshold for Generating a Power Notice	11-14
PoE/PoE+ Allocation Using LLDP Information	11-17
LLDP with PoE	11-17
Displaying the Switch's Global PoE Power Status	11-18
Displaying PoE Status on All Ports	11-20
Displaying the PoE Status on Specific Ports	11-22
Planning and Implementing a PoE Configuration	11-24
Power Requirements	11-24
Assigning PoE Ports to VLANs	11-25
Applying Security Features to PoE Configurations	11-25
Assigning Priority Policies to PoE Traffic	11-25

PoE Event Log Messages	11-27
“Informational” PoE Event-Log Messages	11-27
“Warning” PoE Event-Log Messages	11-28

Introduction to PoE

PoE technology allows IP telephones, wireless LAN access points, and other appliances to receive power and transfer data over existing ethernet LAN cabling. For more information about PoE technology, refer to the *PoE Planning and Implementation Guide*, which is available on the ProCurve Networking web site at www.procurve.com. Select **Support** and then click on **Manuals**.

PoE Terminology

PoE and PoE+ operate similarly in most cases. The CLI commands are the same for a PoE module or a PoE+ zl module. Any differences between PoE and PoE+ operation will be noted, otherwise the term “PoE” is used to designate both PoE and PoE+ functionality.

Term	Use in this Manual
active PoE port	A PoE-enabled port connected to a PD requesting power.
DTE	Data Terminal Equipment
MPS	Maintenance Power Signature; the signal a PD sends to the switch to indicate that the PD is connected and requires power.
Oversubscribed	The state where there are more PDs requesting PoE power than can be accommodated.
PD	Powered Device. This is an IEEE 802.3af-compliant or IEEE 802.3at-compliant device that receives its power through a direct connection to a PoE port in a PoE device. Examples of PDs include Voice-over-IP (VoIP) telephones, wireless access points, and remote video cameras.
PoE	Power-Over-Ethernet; the method by which PDs receive power from a PoE module (operates according to the IEEE 802.3af standard). Some pre-standard PoE devices are also supported; refer to the FAQs for your switch model.
PoE+ (POEP)	Power-over-Ethernet Plus; the method by which PDs receive power according to the 802.3at standard. It is backward compatible with devices using the 802.3af standard.
PoE Module	Refers to a PoE Module for the switches covered in this guide.

Term	Use in this Manual
port-number priority	Refers to the type of power prioritization where, within a priority class, a PoE module assigns the highest priority to the lowest-numbered port in the module, the second-highest priority to the second lowest-numbered port in the module, and so on. Note that power priority rules apply only if PoE provisioning on the module becomes oversubscribed.
priority class	Refers to the type of power prioritization that uses Low (the default), High , and Critical priority assignments to determine which groups of ports will receive power. Note that power priority rules apply only if PoE provisioning becomes oversubscribed.
PSE	Power-Sourcing Equipment. A PSE provides power to IEEE 802.3af-compliant or IEEE 802.3at-compliant PDs directly connected to the ports on the module. The PoE module is an <i>endpoint</i> PSE.

PoE Operation

Using the commands described in this chapter, you can:

- Enable or disable PoE operation on individual ports.
- Monitor PoE status and performance per module.
- Configure a non-default power threshold for SNMP and Event Log reporting of PoE consumption on either all PoE ports on the switch or on all PoE ports in one or more PoE modules.
- Specify the port priority you want to use for provisioning PoE power in the event that the PoE resources become oversubscribed.

A PSE detects the power needed by a PD before supplying that power, a detection phase referred to as “searching”. If the PSE can’t supply the required amount of power, it does not supply any power. For PoE using a Type 1 device, a PSE will not supply any power to a PD unless the PSE has at least 17 W available. For example, if a PSE has a maximum available power of 382W and is already supplying 378W, and is then connected to a PD requiring 10W, the PSE will not supply power to the PD.

For PoE+ using Type 2 devices, the PSE must have at least 33 W available. A slot in a zl chassis can provide a maximum of 370 watts of PoE/PoE+ power to a module.

Configuration Options

In the default configuration, all ports on the PoE module in a ProCurve switch covered in this guide are configured to support PoE operation. You can:

- Disable or re-enable per-port PoE operation on individual ports to help control power usage and avoid oversubscribing PoE resources.
- Configure per-port priority for allocating power in case a PoE module becomes oversubscribed and must drop power for some lower-priority ports to support the demand on other, higher-priority ports.
- Manually allocate the amount of PoE power for a port by usage, value, or class.
- Allocate PoE power based on the link-partner’s capabilities via LLDP.

Note

The ports support standard networking links and PoE links. You can connect either a non-PoE device or a PD to a port enabled for PoE without reconfiguring the port.

PD Support

To best utilize the allocated PoE power, spread your connected PoE devices as evenly as possible across modules. Depending on the amount of power the power supply device delivers to a PoE module, there may or may not always be enough power available to connect and support PoE operation on all the ports in the module. When a new PD connects to a PoE module and the module does not have enough power left for that port:

- If the new PD connects to a port “X” having a *higher* PoE priority than another port “Y” that is already supporting another PD, then the power is removed from port “Y” and delivered to port “X”. In this case the PD on port “Y” loses power and the PD on port “X” receives power.
- If the new PD connects to a port “X” having a *lower* priority than all other PoE ports currently providing power to PDs, then power is not supplied to port “X” until one or more PDs using higher priority ports are removed.

In the default configuration (**usage**), when a PD connects to a PoE port and begins operating, the port retains only enough PoE power to support the PD’s operation. Unused power becomes available for supporting other PD connections. However, if you configure the **poe-allocate-by** option to either **value** or **class**, then all of the power configured is allocated to the port.

For PoE (not PoE+), while 17 watts must be available for a PoE module on the switch to begin supplying power to a port with a PD connected, 17 watts per port is not continually required if the connected PD requires less power. For example, with 20 watts of PoE power remaining available on a module, you can connect one new PD without losing power to any currently connected PDs on that module. If that PD draws only 3 watts, then 17 watts remain available and you can connect at least one more PD to that module without interrupting power to any other PoE devices connected to the same module. If the next PD you connect draws 5 watts, then only 12 watts remain unused. With only 12 unused watts available, if you then connect yet another PD to a higher-priority PoE port, then the lowest-priority port on the module loses PoE power and remains unpowered until the module once again has 17 or more watts available. (For information on power priority, refer to “Power Priority Operation” on page 11-7.)

For PoE+, there must be 33 watts available for the module to begin supplying power to a port with a PD connected. A slot in a zl chassis can provide a maximum of 370 watts of PoE/PoE+ power to a module.

Disconnecting a PD from a PoE port causes the module to stop providing PoE power to that port and makes the power available to any other PoE ports that have PDs connected and waiting for power. If the PD demand for power becomes greater than the PoE power available, then power is transferred from the lower-priority ports to the higher-priority ports. (Ports not currently providing power to PDs are not affected.)

Power Priority Operation

When Is Power Allocation Prioritized?

If a PSE can provide power for all connected PD demand, it does not use its power priority settings to allocate power. However, if the PD power demand oversubscribes the available power, then the power allocation is prioritized to the ports that present a PD power demand. This causes the loss of power from one or more lower-priority ports to meet the power demand on other, higher-priority ports. This operation occurs regardless of the order in which PDs connect to the module's PoE-enabled ports.

How Is Power Allocation Prioritized?

There are two ways that PoE power is prioritized:

- Using a *priority class* method, a power priority of **Low** (the default), **High**, or **Critical** is assigned to each enabled PoE port.
- Using a *port-number priority* method, a lower-numbered port has priority over a higher-numbered port within the same configured priority class, for example, port A1 has priority over port A5 if both are configured with **High** priority.

Configuring PoE Operation

In the default configuration, PoE support is enabled on the ports in a PoE module installed on the switch. The default priority for all ports is **Low** and the default power notification threshold is **80** (%).

Using the CLI, you can:

- Disable or re-enable PoE operation on individual PoE ports
- Enable support for pre-standard devices
- Change the PoE priority level on individual PoE ports
- Change the threshold for generating a power level notice
- Manually allocate the amount of PoE power for a port by usage, value, or class.
- Allocate PoE power based on the link-partner's capabilities via LLDP.

Disabling or Re-Enabling PoE Port Operation

Syntax: [no] interface <port-list> power-over-ethernet

*Re-enables PoE operation on <port-list> and restores the priority setting in effect when PoE was disabled on <port-list>. The **no** form of the command disables PoE operation on <port-list>. (Default: All PoE ports are initially enabled for PoE operation at **Low** priority. If you configure a higher priority, this priority is retained until you change it.)*

Note: For PoE, disabling all ports allows the 22 W of minimum PoE power or 38 W for PoE+ power allocated for the module to be recovered and used elsewhere. You must disable ALL ports for this to occur.

Enabling Support for Pre-Standard Devices

The ProCurve switches covered in this guide also support some pre-802.3af devices. For a list of the devices supported, refer to the FAQs for your switch model.

Syntax: [no] power-over-ethernet pre-std-detect

Detects and powers pre-802.3af standard devices.

Note: This is enabled by default.

Configuring the PoE Port Priority Level

Syntax: interface < port-list > power-over-ethernet [critical | high | low]

Reconfigures the PoE priority level on <port-list>. For a given level, ports are prioritized by port number in ascending order. For example, if ports A1-A24 have a priority level of critical, port A1 has priority over ports A2-A24.

If there is not enough power available to provision all active PoE ports at a given priority level, then the lowest-numbered port at that level will be provisioned first. For chassis switches, the lowest-numbered port at that level starting with module A, then B, C, and so on is provisioned. PoE priorities are invoked only when all active PoE ports cannot be provisioned (supplied with PoE power).

- **Critical:** Specifies the highest-priority PoE support for <port-list>. The active PoE ports at this level are provisioned before the PoE ports at any other level are provisioned.
- **High:** Specifies the second priority PoE support for <port-list>. The active PoE ports at this level are provisioned before the Low priority PoE ports are provisioned.
- **Low:** (the default): Specifies the third priority PoE support for <port-list>. The active PoE ports at this level are provisioned only if there is power available after provisioning any active PoE ports at the higher priority levels.

*In chassis switches, you can use one command to set the same priority level on PoE ports in multiple modules. For example, to configure the priority to **High** for ports c5-c10, C23-C24, D1-D10, and D12, you could use this command:*

```
ProCurve(config)# interface c5-c10,c23-c24,  
d1-d10,d12 power-over-ethernet high
```

Suppose, for example, that you configure the PoE priority for a module in slot C as shown in table 11-1.

Table 11-1. Example of PoE Priority Operation on a PoE Module

Port	Priority Setting	Configuration Command ¹ and Resulting Operation with PDs connected to Ports C3 Through C24
C3 - C17	Critical	<p>In this example, the following CLI command sets ports C3-C17 to Critical:</p> <pre data-bbox="411 348 1182 395">ProCurve(config)# interface c3-c17 power-over-ethernet critical</pre> <p>The Critical priority class always receives power. If there is not enough power to provision PDs on all of the ports configured for this class, then no power goes to ports configured for High and Low priority. If there is enough power to provision PDs on only some of the critical-priority ports, then power is allocated to these ports in ascending order, beginning with the lowest-numbered port in the class, which, in this case, is port 3.</p>
C18 - C21	High	<p>In this example, the following CLI command sets ports C19-C22 to High:</p> <pre data-bbox="411 644 1270 670">ProCurve(config)# interface c19-c22 power-over-ethernet high</pre> <p>The High priority class receives power only if all PDs on ports with a Critical priority setting are receiving power. If there is not enough power to provision PDs on all ports with a high priority, then no power goes to ports with a low priority. If there is enough power to provision PDs on only some of the high-priority ports, then power is allocated to these ports in ascending order, beginning, in this example, with port 18, until all available power is in use.</p>
C22 - C24	Low	<p>In this example, the CLI command sets ports C23-C24 to Low²:</p> <pre data-bbox="411 911 1255 937">ProCurve(config)# interface c23-c24 power-over-ethernet low</pre> <p>This priority class receives power only if all PDs on ports with High and Critical priority settings are receiving power. If there is enough power to provision PDs on only some low-priority ports, then power is allocated to the ports in ascending order, beginning with the lowest-numbered port in the class (port 22, in this case), until all available power is in use.</p>
C1 - C2	- n/a -	<p>In this example, the CLI command disables PoE power on ports C1-C2:</p> <pre data-bbox="411 1161 1210 1187">ProCurve(config)# no interface c1-c2 power-over-ethernet</pre> <p>There is no priority setting for the ports in this example.</p>

¹ For a listing of PoE configuration commands, with descriptions, refer to “Configuring PoE Operation” on page 11-8.

² In the default PoE configuration, the ports are already set to the **low** priority. In this case, the command is not necessary.

PoE Priority With Two or More Modules

Ports across two or more modules can be assigned a class priority of either **Low** (the default), **High**, or **Critical**, for example, A5, B7, and C10 could all be assigned a priority class of **Critical**. When power is allocated to the ports on a priority basis, the **Critical** priority power requests are allocated to module A first, then Module B, C, and so on. Next, the **High** priority power requests are

allocated starting with module A, then B, C, and the remaining modules in order. Any remaining power is allocated in the same manner for the **Low** priority ports, beginning with module A though the remaining modules. If there is not enough PoE power for all the PDs connected to PoE modules in the switch, power is allocated according to priority class across modules. For example:

All ports on module C are prioritized as **Critical**.

```
ProCurve(config)# interface c1-c24 power-over-ethernet
                    critical
```

All ports on module A are prioritized as **Low**.

```
ProCurve(config)# interface a1-a24 power-over-ethernet
                    low
```

There are 48 PDs attached to all ports of modules A and C (24 ports each module).

There is only enough PoE power for 32 ports (8.5 watts x 32 ports = 273 watts).

The result is that all the **Critical** priority ports on module C would receive power, but only 8 ports on module A would receive power.

On module A, the port A1 has the highest priority of the ports in that module if all ports are in the same priority class, which is the case for this example. Since a minimum 17 + 5 watts of power is allocated per PoE module for PoE, port A1 will always receive PoE power. If another port on module A had a higher priority class than port A1, that port would be allocated the power before port A1.

For PoE+ modules there must be a minimum of 33 + 5 watts of power allocated per PoE+ module.

Controlling PoE Allocation

The default option for PoE allocation is **usage**, which is what a PD attached to the port is allocated. You can override this value by specifying the amount of power allocated to a port by using the **class** or **value** options.

Syntax: [no] int <port-list> poe-allocate-by [usage | class | value]

Allows you to manually allocate the amount of PoE power for a port by either its class or a defined value.

usage: *The automatic allocation by a PD*

class: *Uses the power ramp-up signature of the PD to identify which power class the device will be in. Classes and their ranges are shown in table 11-2.*

value: *A user-defined level of PoE power allocated for that port.*

Note

The allowable PD requirements are lower than those specified for PSEs to allow for power losses along the Cat-5 cable.

Table 11-2. Power Classes and Their Values

Power Class	Value
0	Depends on cable type and PoE architecture. Maximum power level output of 15.4 watts at the PSE. This is the default class; if there isn't enough information about the load for a specific classification, the PSE classifies the load as class 0 (zero).
1	Requires at least 4 watts at the PSE.
2	Requires at least 7 watts at the PSE.
3	15.4 watts
4	For PoE+ Maximum power level output of 30 watts at the PSE.

For example, to allocate by class for ports 6 - 8:

```
ProCurve(config)# int 6-8 PoE-allocate-by class
```

Manually Configuring PoE Power Levels

You can specify a power level (in watts) allocated for a port by using the **value** option. This is the maximum amount of power that will be delivered.

To configure a port by value, first set the PoE allocation by entering the **poe-allocate-by value** command:

```
ProCurve(config)# int A6 poe-allocate-by value
```

or in interface context:

```
ProCurve(eth-A6)# poe-allocate-by value
```

Then select a value:

```
ProCurve(config)# int A6 poe-value 15
```

or in interface context:

```
ProCurve(eth-A6)# poe-value 15
```

To view the settings, enter the **show power-over-ethernet** command:

```
ProCurve(config)# show power-over-ethernet A6

Status and Counters - Port Power Status for port A6

Power Enable      : Yes
Priority          : low
AllocateBy       : value
Detection Status : Delivering
LLDP Detect       : enabled
Configured Type  : 15 W
Value            : 15 W
Power Class      : 2
Over Current Cnt : 0
Power Denied Cnt : 0
MPS Absent Cnt  : 0
Short Cnt        : 0
Voltage          : 55.1 V
Current          : 154 mA
```

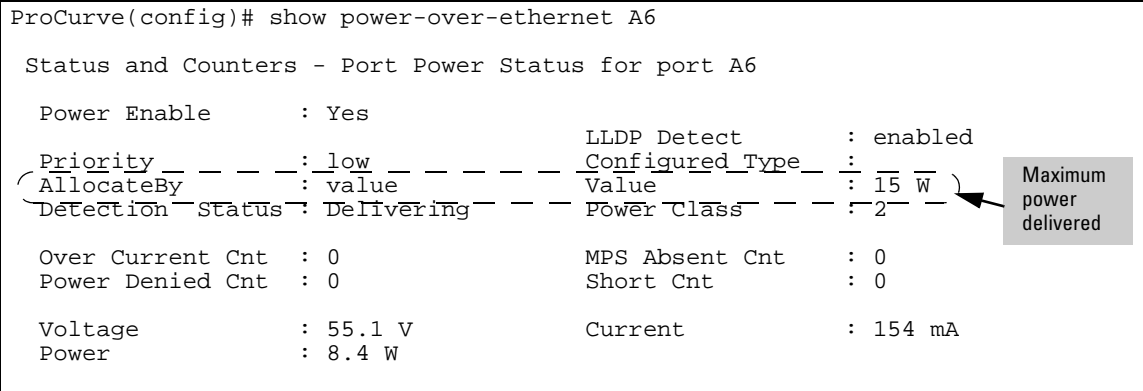


Figure 11-1. Example Displaying PoE Allocation by Value and the Maximum Power Delivered

If you set the PoE maximum value to less than the PD requires, a fault occurs.

```
ProCurve(config)# int A7 poe-value 4

ProCurve(config)# show power-over-ethernet A7

Status and Counters - Port Power Status for port A7

  Power Enable      : Yes
  Priority          : low
  AllocateBy _ _ _ : value _ _ _
  (Detection Status : fault _ _ _ )
  LLDP Detect      : enabled
  Configured Type  :
  Value           : 4 W
  Power Class     : 2

  Over Current Cnt : 1
  Power Denied Cnt : 2
  MPS Absent Cnt  : 0
  Short Cnt       : 0

  Voltage         : 55.1 V
  Power           : 8.4 W
  Current        : 154 mA
```

Figure 11-2. Example Showing PoE Power Value Set Too Low for the PD

Configuring PoE Redundancy (Chassis Switches Only)

When PoE redundancy is enabled, PoE redundancy occurs automatically. The switch keeps track of power use and won't supply PoE power to additional PoE devices trying to connect if that results in the switch not having enough power in reserve for redundancy if one of the power supplies should fail.

Syntax: [no] power-over-ethernet redundancy [n+1 | full]

Allows you to set the amount of power held in reserve for redundancy.

*The **no** option means that all available power can be allocated to PDs.*

Default: No PoE redundancy enforced.

n+1: *One of the power supplies is held in reserve for redundancy. If a single power supply fails, no powered devices are shut down. If power supplies with different ratings are used, the highest-rated power supply is held in reserve to ensure full redundancy.*

full: *Half of the available power supply is held in reserve for redundancy. If power supplies with different ratings are used, the highest-rated power supply is held in reserve to ensure full redundancy.*

See the *PoE Planning and Implementation Guide* for more information about PoE redundancy and power supplies, available on the ProCurve website at www.procurve.com

Changing the Threshold for Generating a Power Notice

You can configure one of the following thresholds:

- A global power threshold that applies to all modules on the switch. This setting acts as a trigger for sending a notice when the PoE power consumption on any PoE module installed in the switch crosses the configured global threshold level. (Crossing the threshold level in either direction—PoE power usage either increasing or decreasing—triggers the notice.) The default setting is 80%.
- A per-slot power threshold that applies to an individual PoE module installed in the designated slot. This setting acts as a trigger for sending a notice when the module in the specified slot exceeds or goes below a specific level of PoE power consumption.

Syntax: power-over-ethernet [slot < slot-id-range >] threshold < 1 - 99 >

This command specifies the PoE usage level (as a percentage of the PoE power available on a module) at which the switch generates a power usage notice. This notice appears as an SNMP trap and a corresponding Event Log message, and occurs when a PoE module's power consumption crosses the configured threshold value. That is, the switch generates a notice whenever the power consumption on a module either exceeds or drops below the specified percentage of the total PoE power available on the module.

This command configures the notification threshold for PoE power usage on either a global or per-module (slot) basis.

Without the [slot <slot-id-range>] option, the switch applies one power threshold setting on all PoE modules installed in the switch. For example, suppose slots A, B, and C each have a PoE module installed. In this case, executing the following command sets the global notification threshold to 70% of available PoE power.

```
ProCurve(config)# power-over-ethernet threshold  
70
```

With this setting, if module B is allocated 100 watts of PoE power and is using 68 watts, and then another PD is connected to the module in slot B that uses 8 watts, the 70% threshold of 70 watts is exceeded. The switch sends an SNMP trap and generates this Event Log message:

Slot B POE usage has exceeded threshold of 70%.

If the switch is configured for debug logging, it also sends the Event Log message to the configured debug destination(s).

On any PoE module, if an increasing PoE power load (1) exceeds the configured power threshold (which triggers the log message and SNMP trap), and then (2) later decreases and drops below the threshold again, the switch generates another SNMP trap, plus a message to the Event Log and any configured Debug destinations.

Syntax: power-over-ethernet [slot <slot-id-range>] threshold <1 - 99 >
(Continued)

To continue the preceding example, if the PoE power usage on the PoE module in slot B drops below 70%, another SNMP trap is generated and you will see this message in the Event Log:

Slot B POE usage is below threshold of 70%.

*For a message listing, refer to “” on page 11-28. (Default Global PoE Power Threshold: **80**). By using the [slot <slot-id-range>] option, you can specify different notification thresholds for different PoE modules installed in the switch. For example, you could set the power threshold for a PoE module in slot “A” to 75% and the threshold for the module in slot “B” to 68% by executing the following two commands:*

```
ProCurve(config)# power-over-ethernet slot a  
threshold 75
```

```
ProCurve(config)# power-over-ethernet slot b  
threshold 68
```

*Note that the last **threshold** command affecting a given slot supersedes the previous threshold command affecting the same slot. Thus, executing the following two commands in the order shown sets the threshold for the PoE module in slot “D” to 75%, but leaves the thresholds for any PoE modules in the other slots at 90%.*

```
ProCurve(config)# power-over-ethernet  
threshold 90
```

```
ProCurve(config)# power-over-ethernet slot d  
threshold 75
```

(If you reverse the order of the above two commands, all PoE modules in the switch will have a threshold of 90%.)

PoE/PoE+ Allocation Using LLDP Information

LLDP with PoE

When using PoE, enabling **poe-lldp-detect** allows automatic power configuration if the link partner supports PoE. When LLDP is enabled, the information about the power usage of the PD is available and the switch can then comply with or ignore this information. You can configure PoE on each port according to the PD (IP phone, wireless device, etc.) specified in the LLDP field. The default configuration is for PoE information to be ignored if detected through LLDP.

Note

Detecting PoE information via LLDP only affects power delivery; it does not affect normal Ethernet connectivity.

To enable or disable ports for allocating power using LLDP, use this command.

Syntax: `int <port-list> poe-lldp-detect [enabled | disabled]`

Enables or disables port(s) for allocating PoE power based on the link-partner's capabilities via LLDP.

Default: Disabled

For example, you can enter this command to enable LLDP detection:

```
ProCurve(config)# int A7 poe-lldp-detect enabled
```

or in interface context:

```
ProCurve(eth-A7)# poe-lldp-detect enabled
```

To enable PoE detection via LLDP TLV advertisement, use this command and insert the desired port or ports:

```
ProCurve(config)# lldp config <port-number>  
medTlvenable poe
```

Displaying the Switch's Global PoE Power Status

Syntax: `show power-over-ethernet [brief | [ethernet] <port-list> | [slot <slot-id-range> | all]]`

Displays the switch's global PoE power status, including:

- **Total Available Power:** *Lists the maximum PoE wattage available to provision active PoE ports on the switch. This is the amount of usable power for PDs.*
- **Total Failover Power:** *Lists the amount of PoE power available in the event of a single power supply failure. This is the amount of power the switch can maintain without dropping any PDs.*
- **Total Redundancy Power:** *Indicates the amount of PoE power that is held in reserve for redundancy in case of a power supply failure.*
- **Total Remaining Power:** *The amount of PoE power still available.*

brief: *Displays PoE information for each port. See “Displaying PoE Status on All Ports” on page 11-21.*

<port-list>: *Displays PoE information for the ports in <port-list>. See “Displaying the PoE Status on Specific Ports” on page 11-23.*

<slot-id-range>: *Displays PoE information for the selected slots. (See figure 11-5). Enter the **all** option to display the PoE information for all slots.*

For example, **show power-over-ethernet** displays data similar to that in figure 11-3.

Power Over Ethernet (PoE/PoE+) Operation
Displaying the Switch's Global PoE Power Status

```
ProCurve(config)# show power-over-ethernet

Status and Counters - System Power Status

Pre-standard Detect      : On
System Power Status     : No redundancy
PoE Power Status        : No redundancy

Chassis power-over-ethernet:

Total Available Power   : 600 W
Total Failover Power    : 300 W
Total Redundancy Power  : 0 W
Total used Power        : 9 W +/- 6W
Total Remaining Power   : 591 W

Internal Power
 1 300W/POE /Connected.
 2 300W/POE /Connected.
 3 Not Connected.
 4 Not Connected.

External Power
EPS1 /Not Connected.
EPS2 /Not Connected.
```

Figure 11-3. Example of show power-over-ethernet Command Output

Displaying PoE Status on All Ports

Syntax: show power-over-ethernet brief

Displays the following port power status:

- **PoE Port:** Lists all PoE-capable ports on the switch.
- **Power Enable:** Shows **Yes** for ports enabled to support PoE (the default) and **No** for ports on which PoE is disabled.
- **Power Priority:** Lists the power priority (**Low**, **High**, and **Critical**) configured on ports enabled for PoE. (For more on this topic, refer to the power command description under “Configuring PoE Operation” on page 11-8.)
- **Alloc by:** Displays how PoE is allocated (**usage, class, value**)
- **Alloc Power:** The maximum amount of PoE power allocated for that port (expressed in watts). Default: 17W for PoE; 33W for PoE+.
- **Actual Power:** The power actually being used on that port.
- **Configured Type:** If configured, shows the user-specified identifier for the port. If not configured, the field is empty.
- **Detection Status:**
 - **Searching:** The port is trying to detect a PD connection.
 - **Delivering:** The port is delivering power to a PD.
 - **Disabled:** On the indicated port, either PoE support is disabled or PoE power is enabled but the PoE module does not have enough power available to supply the port's power needs.
 - **Fault:** The switch detects a problem with the connected PD.
 - **Other Fault:** The switch has detected an internal fault that prevents it from supplying power on that port.
- **Power Class:** Shows the 802.3af power class of the PD detected on the indicated port. Classes include:

0: 0.44w to 12.95w can be drawn by the PD. Default class.	3: 6.49w to 12.95w
1: 0.44w to 3.84w	4: For PoE+; up to 25.5 watts can be drawn by the PD.
2: 3.84w to 6.49w	

Power Over Ethernet (PoE/PoE+) Operation

Displaying the Switch's Global PoE Power Status

For example, **show power-over-ethernet brief** displays this output:

```
ProCurve(config)# show power-over-ethernet brief

Status and Counters - Port Power Status

System Power Status      : No redundancy
PoE Power Status        : No redundancy

Available: 600 W  Used: 9 W  Remaining: 591 W

Module A Power
Available: 408 W  Used: 9 W  Remaining: 399 W

PoE Port | Power Enable | Power Priority | Alloc By | Alloc Power | Actual Power | Configured Type | Detection Status | Power Class
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
A1       | Yes          | low           | usage 17 W | 17 W         | 0.0 W        |                   | Searching        | 0
A2       | Yes          | low           | usage 17 W | 17 W         | 0.0 W        |                   | Searching        | 0
A3       | Yes          | low           | usage 17 W | 17 W         | 0.0 W        |                   | Searching        | 0
A4       | Yes          | low           | usage 17 W | 17 W         | 0.0 W        |                   | Searching        | 0
A5       | Yes          | low           | usage 17 W | 17 W         | 0.0 W        |                   | Searching        | 0
A6       | Yes          | low           | usage 17 W | 17 W         | 8.4 W        |                   | Delivering       | 2
A7       | Yes          | low           | usage 17 W | 17 W         | 0.0 W        |                   | Searching        | 0
A8       | Yes          | low           | usage 17 W | 17 W         | 0.0 W        |                   | Searching        | 0
A9       | Yes          | low           | usage 17 W | 17 W         | 0.0 W        |                   | Searching        | 0
```

Figure 11-4. Example of show power-over-ethernet brief Command Output

You can also show the PoE information by slot:

```
ProCurve(config)# show power-over-ethernet slot A

Status and Counters - System Power Status for slot A

Maximum Power      : 408 W          Operational Status : On
Power In Use       : 9 W +/- 6 W    Usage Threshold (%) : 80
```

Figure 11-5. Showing the PoE Information by Slot

Displaying the PoE Status on Specific Ports

Syntax: show power-over-ethernet <port-list >

Displays the following PoE status and statistics (since the last reboot) for each port in <port-list >:

- **Power Enable:** Shows **Yes** for ports enabled to support PoE (the default) and **No** for ports on which PoE is disabled. Note that for ports on which power is disabled, this is the only field displayed by **show power-over-ethernet < port-list >**.
- **Priority:** Lists the power priority (**Low**, **High**, and **Critical**) configured on ports enabled for PoE. (For more on this topic, refer to the power command description under “Configuring PoE Operation” on page 11-8.)
- **Allocate by:** How PoE is allocated (**usage**, **class**, **value**)
- **Detection Status:**
 - **Searching:** The port is available to support a PD.
 - **Delivering:** The port is delivering power to a PD.
 - **Disabled:** PoE power is enabled on the port but the PoE module does not have enough power available to supply the port's power needs.
 - **Fault:** The switch detects a problem with the connected PD.
 - **Other Fault:** The switch has detected an internal fault that prevents it from supplying power on that port.
- **Over Current Cnt:** Shows the number of times a connected PD has attempted to draw more than 15.4 watts for PoE or 24.5 watts for PoE+. Each occurrence generates an Event Log message.
- **Power Denied Cnt:** Shows the number of times PDs requesting power on the port have been denied due to insufficient power available. Each occurrence generates an Event Log message.
- **Voltage:** The total voltage, in Volts, being delivered to PDs.
- **Power:** The total power, in Watts, being delivered to PDs.
- **LLDP Detect:** Port is enabled or disabled for allocating PoE power based on the link-partner's capabilities via LLDP
- **Configured Type:** If configured, shows the user-specified identifier for the port. If not configured, the field is empty.
- **Value:** The maximum amount of PoE power allocated for that port (expressed in watts). Default: 17W for PoE; 33W for PoE+

Power Over Ethernet (PoE/PoE+) Operation

Displaying the Switch's Global PoE Power Status

- **Power Class:** Shows the power class of the PD detected on the indicated port. Classes include:

0: 0.44w to 12.95w **2:** 3.84w to 6.49w **4:** For PoE+; up to 25.5 watts can be drawn by the PD
1: 0.44w to 3.84w **3:** 6.49w to 12.95w

- **MPS Absent Cnt:** This value shows the number of times a detected PD has no longer requested power from the port. Each occurrence generates an Event Log message. ("MPS" refers to the "Maintenance Power Signature." Refer to "PoE Terminology" on page 11-3.)
- **Short Cnt:** Shows the number of times the switch provided insufficient current to a connected PD.
- **Current:** The total current, in mA, being delivered to PDs.

For example, if you wanted to view the PoE status of ports A6 and A7, you would use **show power-over-ethernet A6-A7** to display the data:

```
ProCurve(config)# show power-over-ethernet A6-A7

Status and Counters - Port Power Status for port A6

Power Enable      : Yes
Priority           : low
AllocateBy        : value
Detection Status  : Delivering
LLDP Detect       : enabled
Configured Type   :
Value             : 17 W
Power Class       : 2

Over Current Cnt  : 0
Power Denied Cnt : 0
MPS Absent Cnt   : 0
Short Cnt        : 0

Voltage           : 55.1 V
Power             : 8.4 W
Current          : 154 mA

Status and Counters - Port Power Status for port A7

Power Enable      : yes
Priority           : low
AllocateBy        : value
Detection Status  : Searching
LLDP Detect       : disabled
Configured Type   :
Value             : 17 W
Power Class       : 0

Over Current Cnt  : 0
Power Denied Cnt : 0
MPS Absent Cnt   : 0
Short Cnt        : 0

Voltage           : 0 V
Power             : 0 W
Current          : 0 mA
```

Figure 11-6. Example of Show Power-Over-Ethernet < port-list > Output

Planning and Implementing a PoE Configuration

This section provides an overview of some considerations for planning a PoE application. For additional information on this topic, refer to the *HP ProCurve PoE Planning and Implementation Guide* which is available on the ProCurve Networking web site at www.procurve.com. Select **Support**, and then click on **Manuals**.

Some of the elements you may want to consider for a PoE installation include:

- Port assignments to VLANs
- Use of security features
- Power requirements

This section can help you to plan your PoE installation. If you use multiple VLANs in your network, or if you have concerns about network security, you should read the first two topics. If your PoE installation comes close to (or is likely to exceed) the system's ability to supply power to all devices that may request it, then you should also read the third topic. (If it is unlikely that your installation will even approach a full utilization of the PoE power available, then you may find it unnecessary to spend much time on calculating PoE power scenarios.)

Power Requirements

In order to get the best PoE performance, you should provide enough PoE power to exceed the maximum amount of power that is needed by all the PDs that are being used.

By connecting an external power supply you can optionally provision more PoE wattage per port and or supply the switch with redundant 12V power to operate should an internal power supply fail.

By installing a second power supply in the 5406zl/8206zl or a third power supply in a 5412zl/8212zl chassis, depending on how many PoE ports are being supplied with power, the switch can have redundant power if one power supply fails. A Power Supply Shelf (external power supply) can also be connected to the 5400zl/8200zl switches to provide extra or redundant PoE power.

For example, if the 5406zl has two 24-port PoE modules (J8702A) installed, and all ports are using 15.4 watts, then the total wattage used is 739.2 watts (48 x 15.4). To supply the necessary PoE wattage a J8713A power supply is installed in one of the power supply slots.

To gain redundant power, a second J8713A must be installed in the second power supply slot. If the first power supply fails, then the second power supply can supply all necessary power.

See the *HP ProCurve PoE Planning and Implementation Guide* for detailed information about the PoE/PoE+ power requirements for your switch.

Assigning PoE Ports to VLANs

If your network includes VLANs, you may want to assign various PoE-configured ports to specific VLANs. For example, if you are using PoE telephones in your network, you may want to assign ports used for telephone access to a VLAN reserved for telephone traffic.

Applying Security Features to PoE Configurations

You can utilize security features built into the switch to control device or user access to the network through PoE ports in the same way as non-PoE ports.

- **MAC Address Security:** Using Port Security, you can configure each switch port with a unique list of MAC addresses for devices that are authorized to access the network through that port. For more information, refer to the chapter titled “Configuring and Monitoring Port Security” in the *Access Security Guide* for your switch.

Assigning Priority Policies to PoE Traffic

You can use the configurable QoS (Quality of Service) features in the switch to create prioritization policies for traffic moving through PoE ports. Table 11-3 lists the available classifiers and their order of precedence.

Table 11-3. Classifiers for Prioritizing Outbound Packets

Priority	QoS Classifier
1	UDP/TCP Application Type (port)
2	Device Priority (destination or source IP address)
3	IP Type of Service (ToS) field (IP packets only)
4	VLAN Priority
5	Incoming source-port on the switch
6	Incoming 802.1p priority (present in tagged VLAN environments)

For more on this topic, refer to the chapter titled “Quality of Service: Managing Bandwidth More Effectively” in the *Advanced Traffic Management Guide* for your switch.

PoE Event Log Messages

PoE operation generates these Event Log messages. You can also configure the switch to send these messages to a configured debug destination (terminal device or SyslogD server).

“Informational” PoE Event-Log Messages

Message	Meaning
I <MM/DD/YY> <HH:MM:SS> <chassis ports>	Message header, with severity, date, system time, and system module type (chassis or ports). For more information on Event Log operation, including severity indicators, refer to “Using the Event Log for Troubleshooting Switch Problems” on page C-27
Slot <slot-id> POE usage is below configured threshold of <1-99>%	Indicates that POE usage on the module in the indicated slot has decreased below the threshold specified by the last execution of the power threshold command affecting that module. This message occurs if, after the last reboot, the PoE demand on the module exceeded the power threshold and then later dropped below the threshold value.
port <port-id> applying power to PD	A PoE device is connected to the indicated port and receiving power.
port <port-id> PD detected	The switch has detected a PoE device connected to the indicated port.
Slot <slot-id> software update started on PoE controller <controller-id>	A module needs to have its PoE firmware updated and the software begins the update process. On ProCurve 8212zl switches the controller-id is always “1”
Slot <slot-id> software update completed on PoE controller <controller-id>	A module has its PoE firmware updated and the software has finished this process.

“Warning” PoE Event-Log Messages

Message	Meaning
W <MM/DD/YY> <HH:MM:SS> chassis	Message header, with severity, date, system time, and system module type. For more information on Event Log operation, including severity indicators, refer to “Using the Event Log for Troubleshooting Switch Problems” on page C-27”.
Slot <slot-id> POE usage has exceeded threshold of <1-99>%	Indicates that POE usage in the indicated slot has exceeded the configured threshold for the module, as specified by the last execution of the power threshold or power slot < slot-id > threshold command. (Note that the switch also generates an SNMP trap for this event.)
Port <port-id> PD Denied power due to insufficient power allocation.	There is insufficient power available to power the PD on the indicated port and the port does not have sufficient PoE priority to take power from another active PoE port.
Port <port-id> PD Invalid Signature indication	The switch has detected a non-802.3af-compliant or non-802.3at-compliant device on the indicated port. This message appears for all non-802.3af devices connected to the port, such as other switches, PC-NICs, etc.
Port <port-id> PD MPS Absent indication	The switch no longer detects a device on < port-id >. The device may have been disconnected, powered down, or stopped functioning.
Port <port-id> PD Other Fault indication	There is a problem with the PSE connected to the port.
Port <port-id> PD Over Current indication	The PD connected to < port-id > has requested more than 15.4 watts of power. This may indicate a short-circuit or other problem in the PD.
50v Power Supply is faulted. Failures:<num-failures>	Internal power supply has faulted.
50v Power Supply is OK. Failures: <num-failures>	Internal power supply is now OK.
FET bad on port <port-id>	External FET (Field Effect Transistor) on the port has gone bad and cannot deliver power.

Power Over Ethernet (PoE/PoE+) Operation
PoE Event Log Messages

Port Trunking

Contents

Overview	12-3
Port Trunk Features and Operation	12-5
Trunk Configuration Methods	12-6
Menu: Viewing and Configuring a Static Trunk Group	12-10
CLI: Viewing and Configuring Port Trunk Groups	12-12
Using the CLI To View Port Trunks	12-12
Using the CLI To Configure a Static or Dynamic Trunk Group ...	12-15
Web: Viewing Existing Port Trunk Groups	12-18
Trunk Group Operation Using LACP	12-19
Default Port Operation	12-22
LACP Notes and Restrictions	12-23
Distributed Trunking	12-27
Overview	12-27
Distributed Trunking Interconnect Protocol (DTIP)	12-29
Configuring Distributed Trunking	12-30
ISC Port Configuration	12-30
Distributed Trunking Port Configuration	12-30
Displaying Distributed Trunking Information	12-31
Maximum DT Trunks and Links Supported	12-32
Forwarding Traffic with Distributed Trunking and Spanning Tree	12-32
Forwarding Unicast Traffic Upstream	12-32
Forwarding Broadcast, Multicast, and Unknown Traffic Upstream	12-33
Forwarding Unicast Traffic Downstream (to the Server)	12-33

Forwarding Broadcast, Multicast, and Unknown Traffic Downstream (to the Server)	12-33
Distributed Trunking Restrictions	12-35
Trunk Group Operation Using the “Trunk” Option	12-36
How the Switch Lists Trunk Data	12-37
Outbound Traffic Distribution Across Trunked Links	12-37

Overview

This chapter describes creating and modifying port trunk groups. This includes non-protocol trunks and LACP (802.3ad) trunks.

Port Status and Configuration Features

Feature	Default	Menu	CLI	Web
viewing port trunks	n/a	page 12-10	page 12-12	page 12-18
configuring a static trunk group	none	page 12-10	page 12-16	—
configuring a dynamic LACP trunk group	disabled	—	page 12-16	—

Port trunking allows you to assign up to eight physical links to one logical link (trunk) that functions as a single, higher-speed link providing dramatically increased bandwidth. This capability applies to connections between backbone devices as well as to connections in other network areas where traffic bottlenecks exist. A *trunk group* is a set of up to eight ports configured as members of the same port trunk. Note that the ports in a trunk group do not have to be consecutive. For example:

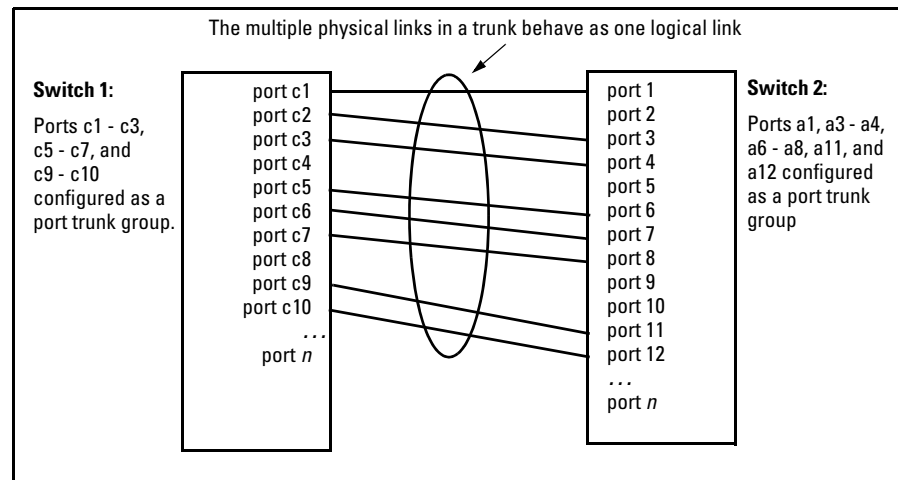


Figure 12-1. Conceptual Example of Port Trunking

With full-duplex operation in a eight-port trunk group, trunking enables the following bandwidth capabilities:

Port Connections and Configuration: All port trunk links must be point-to-point connections between a switch and another switch, router, server, or workstation configured for port trunking. No intervening, non-trunking devices are allowed. It is important to note that ports on both ends of a port trunk group must have the same mode (speed and duplex) and flow control settings.

Note

Link Connections. The switch does not support port trunking through an intermediate, non-trunking device such as a hub, or using more than one media type in a port trunk group. Similarly, for proper trunk operation, all links in the same trunk group must have the same speed, duplex, and flow control.

Port Security Restriction. Port security does not operate on a trunk group. If you configure port security on one or more ports that are later added to a trunk group, the switch resets the port security parameters for those ports to the factory-default configuration.

Caution

To avoid broadcast storms or loops in your network while configuring a trunk, first disable or disconnect all ports you want to add to or remove from the trunk. After you finish configuring the trunk, enable or re-connect the ports.

Port Trunk Features and Operation

The switches covered in this guide offer these options for port trunking:

- LACP: IEEE 802.3ad—page 12-19
- Trunk: Non-Protocol—page 12-36

Up to 144 trunk groups are supported on the switches covered in this guide. The actual maximum depends on the number of ports available on the switch and the number of links in each trunk. (Using the Link Aggregation Control Protocol—LACP—option, you can include standby trunked ports in addition to the maximum of eight actively trunking ports.) The trunks do not have to be the same size, for example, 100 two-port trunks and 11 eight-port trunks are supported.

LACP Note

LACP requires full-duplex (FDx) links of the same media type (10/100Base-T, 100FX, etc.) and the same speed, and enforces speed and duplex conformance across a trunk group. For most installations, ProCurve recommends that you leave the port Mode settings at **Auto** (the default). LACP also operates with **Auto-10**, **Auto-100**, and **Auto-1000** (if negotiation selects FDx), and **10FDx**, **100FDx**, and **1000FDx** settings. (The 10-gigabit ports available for some switch models allow only the **Auto** setting.)

Fault Tolerance: If a link in a port trunk fails, the switch redistributes traffic originally destined for that link to the remaining links in the trunk. The trunk remains operable as long as there is at least one link in operation. If a link is restored, that link is automatically included in the traffic distribution again. The LACP option also offers a standby link capability, which enables you to keep links in reserve for service if one or more of the original active links fails. Refer to “Trunk Group Operation Using LACP” on page 12-19.)

Trunk Configuration Methods

Dynamic LACP Trunk: The switch automatically negotiates trunked links between LACP-configured ports on separate devices, and offers one dynamic trunk option: LACP. To configure the switch to initiate a dynamic LACP trunk with another device, use the **interface** command in the CLI to set the default LACP option to **Active** on the ports you want to use for the trunk. For example, the following command sets ports C1-C4 to LACP active:

```
ProCurve(config) int c1-c4 lacp active
```

Note that the preceding example works if the ports are not already operating in a trunk. To change the LACP option on ports already operating as a trunk, you must first remove them from the trunk. For example, if ports C1 - C4 were LACP-active and operating in a trunk with another device, you would do the following to change them to LACP-passive:

```
ProCurve(config)# no int c1-c4 lacp
```

Removes the ports from the trunk.

```
ProCurve(config)# int c1-c4 lacp passive
```

Configures LACP passive.

Static Trunk: The switch uses the links you configure with the Port/Trunk Settings screen in the menu interface or the **trunk** command in the CLI to create a static port trunk. The switch offers two types of static trunks: LACP and Trunk.

Table 12-1. Trunk Types Used in Static and Dynamic Trunk Groups

Trunking Method	LACP	Trunk
Dynamic	Yes	No
Static	Yes	Yes

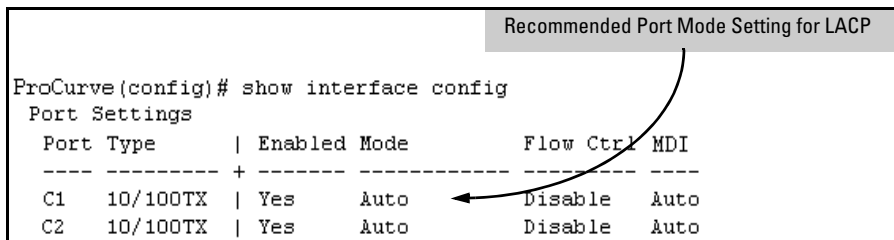
Table 12-2. Trunk Configuration Protocols

Protocol	Trunking Options
LACP (802.3ad)	<p>Provides dynamic and static LACP trunking options.</p> <ul style="list-style-type: none"> • Dynamic LACP — Use the switch-negotiated dynamic LACP trunk when: <ul style="list-style-type: none"> – The port on the other end of the trunk link is configured for Active or Passive LACP. – You want fault-tolerance for high-availability applications. If you use an eight-link trunk you can also configure one or more additional links to operate as standby links that will activate only if another active link goes down. • Static LACP — Use the manually configured static LACP trunk when: <ul style="list-style-type: none"> – The port on the other end of the trunk link is configured for a static LACP trunk – You want to configure non-default spanning tree or IGMP parameters on an LACP trunk group. – <i>You want an LACP trunk group to operate in a VLAN other than the default VLAN and GVRP is disabled. (Refer to “VLANs and Dynamic LACP” on page 12-24.)</i> – You want to use a monitor port on the switch to monitor an LACP trunk. <p>For more information, refer to “Trunk Group Operation Using LACP” on page 12-19.</p>
Trunk (non- protocol)	<p>Provides manually configured, static-only trunking to:</p> <ul style="list-style-type: none"> • Most ProCurve switches and routing switches not running the 802.3ad LACP protocol. • Windows NT and HP-UX workstations and servers <p>Use the Trunk option when:</p> <ul style="list-style-type: none"> – The device to which you want to create a trunk link is using a non-802.3ad trunking protocol – You are unsure which type of trunk to use, or the device to which you want to create a trunk link is using an unknown trunking protocol. – You want to use a monitor port on the switch to monitor traffic on a trunk. <p>Refer to “” on page 12-36.</p>

Table 12-3. General Operating Rules for Port Trunks

Media: For proper trunk operation, all ports on both ends of a trunk group must have the same media type and mode (speed and duplex). (For the switches covered in this guide, ProCurve recommends leaving the port Mode setting at **Auto** or, in networks using Cat 3 cabling, **Auto-10**.)

Port Configuration: The default port configuration is **Auto**, which enables a port to sense speed and negotiate duplex with an Auto-Enabled port on another device. ProCurve recommends that you use the **Auto** setting for all ports you plan to use for trunking. Otherwise, you must manually ensure that the mode setting for each port in a trunk is compatible with the other ports in the trunk.



```
ProCurve(config)# show interface config
Port Settings
-----+-----
C1  10/100TX | Yes  Auto  Disable  Auto
C2  10/100TX | Yes  Auto  Disable  Auto
```

Recommended Port Mode Setting for LACP

Figure 12-2. Recommended Port Mode Setting for LACP

All of the following operate on a per-port basis, regardless of trunk membership:

- Enable/Disable
- Flow control (Flow Ctrl)

LACP is a full-duplex protocol. Refer to “Trunk Group Operation Using LACP” on page 12-19.

Trunk Configuration: All ports in the same trunk group must be the same trunk type (LACP or Trunk). All LACP ports in the same trunk group must be either all static LACP or all dynamic LACP.

A trunk appears as a single port labeled **Dyn1** (for an LACP dynamic trunk) or **Trk1** (for a static trunk of type: LACP, Trunk) on various menu and CLI screens. For a listing of which screens show which trunk types, refer to “How the Switch Lists Trunk Data” on page 12-37.

For spanning-tree or VLAN operation, configuration for all ports in a trunk is done at the trunk level. (You cannot separately configure individual ports within a trunk for spanning-tree or VLAN operation.)

Traffic Distribution: All of the switch trunk protocols use the SA/DA (Source Address/Destination Address) method of distributing traffic across the trunked links. Refer to “Outbound Traffic Distribution Across Trunked Links” on page 12-37.

Spanning Tree: 802.1D (STP) and 802.1w (RSTP) Spanning Tree operate as a global setting on the switch (with one instance of Spanning Tree per switch). 802.1s (MSTP) Spanning Tree operates on a per-instance basis (with multiple instances allowed per switch). For each Spanning Tree instance, you can adjust Spanning Tree parameters on a per-port basis. A static trunk of any type appears in the Spanning Tree configuration display, and you can configure Spanning Tree parameters for a static trunk in the same way that you would configure Spanning Tree parameters on a non-trunked port. (Note that the switch lists the trunk by name—such as **Trk1**—and does not list the individual ports in the trunk.) For example, if ports C1 and C2 are configured as a static trunk named **Trk1**, they are listed in the Spanning Tree display as **Trk1** and do not appear as individual ports in the Spanning Tree displays.

In this example showing part of the show spanning-tree listing, ports C1 and C2 are members of TRK1 and do not appear as individual ports in the port configuration part of the listing.	Port	Type	Cost	Priority	State	Designated Bridge
	C3	100/1000T	5	128	Forwarding	0020c1-b27ac0
	C4	100/1000T	5	128	Forwarding	0060b0-889e00
	C5	100/1000T	5	128	Disabled	
	C6	100/1000T	5	128	Disabled	
	Trk1		1	64	Forwarding	0001e7-a0ec00

Figure 12-3. Example of a Port Trunk in a Spanning Tree Listing

When Spanning Tree forwards on a trunk, all ports in the trunk will be forwarding. Conversely, when Spanning Tree blocks a trunk, all ports in the trunk are blocked.

Note: A dynamic LACP trunk operates only with the default Spanning Tree settings. Also, this type of trunk appears in the CLI **show spanning-tree** display, but not in the Spanning Tree Operation display of the Menu interface.

If you remove a port from a static trunk, the port retains the same Spanning Tree settings that were configured for the trunk.

IP Multicast Protocol (IGMP): A static trunk of any type appears in the IGMP configuration display, and you can configure IGMP for a static trunk in the same way that you would configure IGMP on a non-trunked port. (Note that the switch lists the trunk by name—such as **Trk1**—and does not list the individual ports in the trunk.) Also, creating a new trunk automatically places the trunk in IGMP Auto status if IGMP is enabled for the default VLAN. A dynamic LACP trunk operates only with the default IGMP settings and does not appear in the IGMP configuration display or **show ip igmp** listing.

VLANs: Creating a new trunk automatically places the trunk in the DEFAULT_VLAN, regardless of whether the ports in the trunk were in another VLAN. Similarly, removing a port from a trunk group automatically places the port in the default VLAN. You can configure a static trunk in the same way that you configure a port for membership in any VLAN.

Note: For a dynamic LACP trunk to operate in a VLAN other than the default VLAN (DEFAULT_VLAN), GVRP must be enabled. Refer to “Trunk Group Operation Using LACP” on page 12-19.

Port Security: Trunk groups (and their individual ports) cannot be configured for port security, and the switch excludes trunked ports from the **show port-security** listing. If you configure non-default port security settings for a port, then subsequently try to place the port in a trunk, you will see the following message and the command will not be executed:
<port-list> Command cannot operate over a logical port.

Monitor Port:

Note: A trunk cannot be a monitor port. A monitor port can monitor a static trunk but cannot monitor a dynamic LACP trunk.

Menu: Viewing and Configuring a Static Trunk Group

Important

Configure port trunking *before* you connect the trunked links to another switch, routing switch, or server. Otherwise, a broadcast storm could occur. (If you need to connect the ports before configuring them for trunking, you can temporarily disable the ports until the trunk is configured. Refer to “Enabling or Disabling Ports and Configuring Port Mode” on page 10-15.)

To View and/or Configure Static Port Trunking: This procedure uses the Port/Trunk Settings screen to configure a static port trunk group on the switch.

1. Follow the procedures in the Important note above.
2. From the Main Menu, Select:
2. Switch Configuration ...
2. Port/Trunk Settings
3. Press [E] (for **E**dit) and then use the arrow keys to access the port trunk parameters.

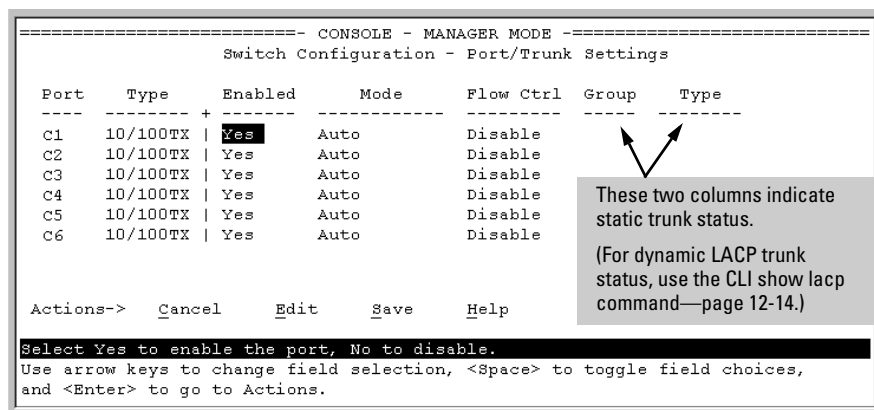


Figure 12-4. Example of the Menu Screen for Configuring a Port Trunk Group

4. In the Group column, move the cursor to the port you want to configure.
5. Use the Space bar to choose a trunk group assignment (**Trk1**, **Trk2**, and so on) for the selected port.

- For proper trunk operation, all ports in a trunk must have the same media type and mode (such as 10/100TX set to 100FDx, or 100FX set to 100FDx). The flow control settings must also be the same for all ports in a given trunk. To verify these settings, refer to “Viewing Port Status and Configuring Port Parameters” on page 10-3.
- You can configure the trunk group with up to eight ports per trunk. If multiple VLANs are configured, all ports within a trunk will be assigned to the same VLAN or set of VLANs. (With the 802.1Q VLAN capability built into the switch, more than one VLAN can be assigned to a trunk. Refer to the chapter titled “Static Virtual LANs (VLANs)” in the *Advanced Traffic Management Guide* for your switch.)

(To return a port to a non-trunk status, keep pressing the Space bar until a blank appears in the highlighted Group value for that port.)

```
===== CONSOLE - MANAGER MODE =====
Switch Configuration - Port/Trunk Settings

Port   Type      Enabled  Mode      Flow Ctrl  Group  Type
-----+-----
C1     10/100TX  | Yes    Auto      Disable    -----
C2     10/100TX  | Yes    Auto      Disable    -----
C3     10/100TX  | Yes    Auto      Disable    -----
C4     10/100TX  | Yes    Auto      Disable    -----
C5     10/100TX  | Yes    Auto      Disable    Trk1  Trunk
C6     10/100TX  | Yes    Auto      Disable    Trk1  Trunk

Actions->  Cancell  Edit    Save    Help

Select whether the port is part of a trunk or Mesh.
Use arrow keys to change field selection, <Space> to toggle field choices,
and <Enter> to go to Actions.
```

Figure 12-5. Example of the Configuration for a Two-Port Trunk Group

6. Move the cursor to the Type column for the selected port and use the Space bar to select the trunk type:
 - LACP
 - Trunk (the default type if you do not specify a type)

All ports in the same trunk group on the same switch must have the same Type (**LACP** or **Trunk**).

7. When you are finished assigning ports to the trunk group, press **[Enter]**, then **[S]** (for **Save**) and return to the Main Menu. (It is not necessary to reboot the switch.)

During the Save process, traffic on the ports configured for trunking will be delayed for several seconds. If the Spanning Tree Protocol is enabled, the delay may be up to 30 seconds.

8. Connect the trunked ports on the switch to the corresponding ports on the opposite device. If you previously disabled any of the trunked ports on the switch, enable them now. (Refer to “Viewing Port Status and Configuring Port Parameters” on page 10-3.)

Check the Event Log (“Using the Event Log for Troubleshooting Switch Problems” on page C-27) to verify that the trunked ports are operating properly.

CLI: Viewing and Configuring Port Trunk Groups

Trunk Status and Configuration Commands

show trunks	below
show lacp	page 12-14
trunk	page 12-16
interface < port-list > lacp	page 12-16

Using the CLI To View Port Trunks

You can list the trunk type and group for all ports on the switch or for selected ports. You can also list LACP-only status information for LACP-configured ports.

Listing Static Trunk Type and Group for All Ports or for Selected Ports.

Syntax: show trunks [< port-list >]

Omitting the < port-list > parameter results in a static trunk data listing for all LAN ports in the switch. For example, in a switch where ports A4 and A5 belong to Trunk 1 and ports A7 and A8 belong to Trunk 2, you have the options shown in figures 12-6 and 12-7 for displaying port data for ports belonging to static trunks.

Using a port list specifies, for switch ports in a static trunk group, only the ports you want to view. In this case, the command specifies ports A5 through A7. However, because port A6 is not in a static trunk group, it does not appear in the resulting listing:

Port A5 appears with an example of a name that you can optionally assign using the Friendly Port Names feature. (Refer to "Using Friendly (Optional) Port Names" on page 10-25.)

```

ProCurve> show trunks e a5-a7

Load Balancing

  Port | Name | Type | Group | Type
-----+-----+-----+-----+-----
  A5   | Print-Server-Trunk | 10/100TX | Trk1 | Trunk
  A7   | not assigned | 10/100TX | Trk2 | Trunk
    
```

Port A6 does not appear in this listing because it is not assigned to a static trunk.

Figure 12-6. Example Listing Specific Ports Belonging to Static Trunks

The **show trunks <port-list>** command in the above example includes a port list, and thus shows trunk group information only for specific ports that have membership in a static trunk. In figure 12-7, the command does not include a port list, so the switch lists all ports having static trunk membership.

```

ProCurve> show trunks

Load Balancing

  Port | Name | Type | Group | Type
-----+-----+-----+-----+-----
  A4   | Print-Server-Trunk | 10/100TX | Trk1 | Trunk
  A5   | Print-Server-Trunk | 10/100TX | Trk1 | Trunk
  A7   | not assigned | 10/100TX | Trk2 | Trunk
  A8   | not assigned | 10/100TX | Trk2 | Trunk
    
```

Figure 12-7. Example of a Show Trunk Listing Without Specifying Ports

Listing Static LACP and Dynamic LACP Trunk Data.

Syntax: show lacp

Lists data for only the LACP-configured ports..

In the following example, ports A1 and A2 have been previously configured for a static LACP trunk. (For more on the “Active” parameter, see table 12-5 on page 12-22.)

```
ProCurve> show lacp
```

LACP					
PORT	LACP	TRUNK	PORT	LACP	LACP
NUMB	ENABLED	GROUP	STATUS	PARTNER	STATUS
----	-----	-----	-----	-----	-----
A1	Active	Trk1	Up	Yes	Success
A2	Active	Trk1	Up	Yes	Success
A3	Active	A3	Down	No	Success
A4	Passive	A4	Down	No	Success
A5	Passive	A5	Down	No	Success
A6	Passive	A6	Down	No	Success

Figure 12-8. Example of a Show LACP Listing

(For a description of each of the above-listed data types, refer to table 12-5, “LACP Port Status Data” on page 12-22.)

Dynamic LACP Standby Links. Dynamic LACP trunking enables you to configure standby links for a trunk by including more than eight ports in a dynamic LACP trunk configuration. When eight ports (trunk links) are up, the remaining link(s) will be held in standby status. If a trunked link that is “Up” fails, it will be replaced by a standby link, which maintains your intended bandwidth for the trunk. (Refer to also the “Standby” entry under “Port Status” in "Table 12-5. LACP Port Status Data" on page 12-22.) In the next example, ports A1 through A9 have been configured for the same LACP trunk. Notice that one of the links shows Standby status, while the remaining eight links are “Up”.

```
ProCurve> show lacp
```

		LACP					
		PORT	LACP	TRUNK	PORT	LACP	LACP
		NUMB	ENABLED	GROUP	STATUS	PARTNER	STATUS
		----	-----	-----	-----	-----	-----
"Up" Links →	A1	Active	Dyn1	Up	Yes	Success	
	A2	Active	Dyn1	Up	Yes	Success	
	A3	Active	Dyn1	Up	Yes	Success	
	A4	Active	Dyn1	Up	Yes	Success	
	A5	Active	Dyn1	Up	Yes	Success	
	A6	Active	Dyn1	Up	Yes	Success	
	A7	Active	Dyn1	Up	Yes	Success	
	A8	Active	Dyn1	Up	Yes	Success	
Standby Link →	A9	Active	Dyn1	Standby	Yes	Success	

Figure 12-9. Example of a Dynamic LACP Trunk with One Standby Link

Using the CLI To Configure a Static or Dynamic Trunk Group

Important

Configure port trunking *before* you connect the trunked links between switches. Otherwise, a broadcast storm could occur. (If you need to connect the ports before configuring them for trunking, you can temporarily disable the ports until the trunk is configured. Refer to “Enabling or Disabling Ports and Configuring Port Mode” on page 10-15.)

The table on page 12-6 describes the maximum number of trunk groups you can configure on the switch. An individual trunk can have up to eight links, with additional standby links if you’re using LACP. You can configure trunk group types as follows:

Trunk Type	Trunk Group Membership	
	TrkX (Static)	DynX (Dynamic)
LACP	Yes	Yes
Trunk	Yes	No

The following examples show how to create different types of trunk groups.

Configuring a Static Trunk or Static LACP Trunk Group.

Syntax: trunk < port-list > < trk1 ... trk144 > < trunk | lacp >

Configures the specified static trunk type.

This example uses ports C4 - C6 to create a non-protocol static trunk group with the group name of **Trk2**.

```
ProCurve(config)# trunk c4-c6 trk2 trunk
```

Removing Ports from a Static Trunk Group. This command removes one or more ports from an existing **Trkx** trunk group.

Caution

Removing a port from a trunk can create a loop and cause a broadcast storm. When you remove a port from a trunk where spanning tree is not in use, ProCurve recommends that you first disable the port or disconnect the link on that port.

Syntax: no trunk < port-list >

Removes the specified ports from an existing trunk group.

For example, to remove ports C4 and C5 from an existing trunk group.

```
ProCurve(config)# no trunk c4-c5
```

Enabling a Dynamic LACP Trunk Group. In the default port configuration, all ports on the switch are set to disabled. To enable the switch to automatically form a trunk group that is dynamic on both ends of the link, the ports on one end of a set of links must be LACP **Active**. The ports on the other end can be either LACP **Active** or LACP **Passive**. The **active** command enables the switch to automatically establish a (dynamic) LACP trunk group when the device on the other end of the link is configured for LACP **Passive**.

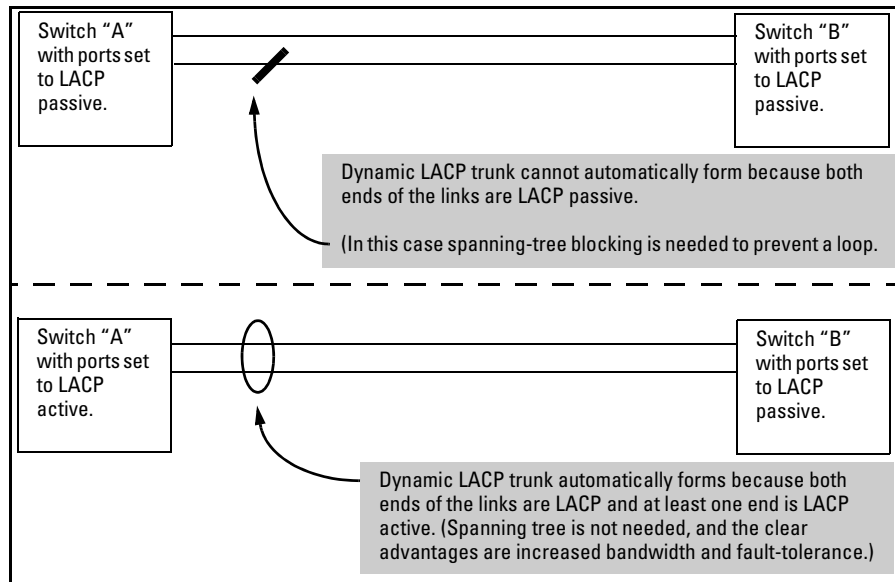


Figure 12-10. Example of Criteria for Automatically Forming a Dynamic LACP Trunk

Syntax: interface <port-list> lacp active

Configures <port-list> as LACP active. If the ports at the other end of the links on <port-list> are configured as LACP passive, then this command enables a dynamic LACP trunk group on <port-list>.

This example uses ports C4 and C5 to enable a dynamic LACP trunk group.

```
ProCurve(config)# interface c4-c5 lacp active
```

Removing Ports from an Dynamic LACP Trunk Group. To remove a port from dynamic LACP trunk operation, you must turn off LACP on the port. (On a port in an operating, dynamic LACP trunk, you cannot change between LACP **Active** and LACP **passive** without first removing LACP operation from the port.)

Caution

Unless spanning tree is running on your network, removing a port from a trunk can result in a loop. To help prevent a broadcast storm when you remove a port from a trunk where spanning tree is not in use, ProCurve recommends that you first disable the port or disconnect the link on that port.

Syntax: no interface <port-list> lacp

Removes <port-list> from any dynamic LACP trunk and returns the ports in <port-list> to passive LACP.

In this example, port C6 belongs to an operating, dynamic LACP trunk. To remove port C6 from the dynamic trunk and return it to passive LACP, you would do the following:

```
ProCurve(config)# no interface c6 lacp
ProCurve(config)# interface c6 lacp passive
```

Note that in the above example, if the port on the other end of the link is configured for active LACP or static LACP, the trunked link will be re-established almost immediately.

Web: Viewing Existing Port Trunk Groups

While the web browser interface does not enable you to configure a port trunk group, it does provide a view of an existing trunk group.

To view any port trunk groups:

Click on the **Status** tab.

Click on **[Port Status]**.

Trunk Group Operation Using LACP

The switch can automatically configure a dynamic LACP trunk group or you can manually configure a static LACP trunk group.

Note

LACP requires full-duplex (FDx) links of the same media type (10/100Base-T, 100FX, etc.) and the same speed, and enforces speed and duplex conformance across a trunk group. For most installations, ProCurve recommends that you leave the port Mode settings at **Auto** (the default). LACP also operates with **Auto-10**, **Auto-100**, and **Auto-1000** (if negotiation selects FDx), and **10FDx**, **100FDx**, and **1000FDx** settings.

LACP trunk status commands include:


Trunk Display Method	Static LACP Trunk	Dynamic LACP Trunk
CLI show lacp command	Included in listing.	Included in listing.
CLI show trunk command	Included in listing.	Not included.
Port/Trunk Settings screen in menu interface	Included in listing.	Not included

Thus, to display a listing of dynamic LACP trunk ports, you must use the **show lacp** command.

In most cases, trunks configured for LACP on the switches covered in this guide operate as described in table 12-4 on the next page.

Table 12-4. LACP Trunk Types

LACP Port Trunk Configuration	Operation
Dynamic LACP	<p>This option automatically establishes an 802.3ad-compliant trunk group, with LACP for the port Type parameter and DynX for the port Group name, where X is an automatically assigned value from 1 to 144, depending on how many dynamic and static trunks are currently on the switch. (The switch allows a maximum of 144 trunk groups in any combination of static and dynamic trunks.)</p> <p>Note: Dynamic LACP trunks operate only in the default VLAN (unless GVRP is enabled and Forbid is used to prevent the trunked ports from joining the default VLAN). Thus, if an LACP dynamic port forms using ports that are not in the default VLAN, the trunk will automatically move to the default VLAN unless GVRP operation is configured to prevent this from occurring. In some cases, this can create a traffic loop in your network. For more on this topic, refer to “VLANs and Dynamic LACP” on page 12-24.</p> <p>Under the following conditions, the switch automatically establishes a dynamic LACP port trunk group and assigns a port Group name:</p> <ul style="list-style-type: none"> • The ports on both ends of each link have compatible mode settings (speed and duplex). • The port on one end of each link must be configured for LACP Active and the port on the other end of the same link must be configured for either LACP Passive or LACP Active. For example:



Either of the above link configurations allow a dynamic LACP trunk link.

Backup Links: A maximum of eight operating links are allowed in the trunk, but, with dynamic LACP, you can configure one or more additional (backup) links that the switch automatically activates if a primary link fails. To configure a link as a standby for an existing eight-port dynamic LACP trunk, ensure that the ports in the standby link are configured as either active-to-active or active-to-passive between switches.

Displaying Dynamic LACP Trunk Data: To list the configuration and status for a dynamic LACP trunk, use the CLI **show lacp** command.

Note: The dynamic trunk is automatically created by the switch, and is not listed in the static trunk listings available in the menu interface or in the CLI **show trunk** listing.

LACP Port Trunk Configuration	Operation
-------------------------------	-----------

Static LACP

Provides a manually configured, static LACP trunk to accommodate these conditions:

- The port on the other end of the trunk link is configured for a static LACP trunk.
- You want to configure non-default spanning tree or IGMP parameters on an LACP trunk group.
- You want an LACP trunk group to operate in a VLAN other than the default VLAN and GVRP is disabled. (Refer to “VLANs and Dynamic LACP” on page 12-24.)
- You want to use a monitor port on the switch to monitor an LACP trunk.

The trunk operates if the trunk group on the opposite device is running one of the following trunking protocols:

- Active LACP
- Passive LACP
- Trunk

This option uses **LACP** for the port Type parameter and **TrkX** for the port Group parameter, where **X** is an automatically assigned value in a range corresponding to the maximum number of trunks the switch allows. (The table on page 12-6 lists the maximum number of trunk groups allowed on the switches covered in this guide.)

Displaying Static LACP Trunk Data: To list the configuration and status for a static LACP trunk, use the CLI **show lacp** command. To list a static LACP trunk with its assigned ports, use the CLI **show trunk** command or display the menu interface Port/Trunk Settings screen.

Static LACP does not allow standby ports.

Default Port Operation

In the default configuration, LACP is disabled for all ports. If LACP is not configured as Active on at least one end of a link, then the port does not try to detect a trunk configuration and operates as a standard, untrunked port. Table 12-5 lists the elements of per-port LACP operation. To display this data for a switch, execute the following command in the CLI:

```
ProCurve> show lacp
```

Table 12-5. LACP Port Status Data

Status Name	Meaning
Port Numb	Shows the physical port number for each port configured for LACP operation (C1, C2, C3 .). Unlisted port numbers indicate that the missing ports are assigned to a static Trunk group are not configured for any trunking.
LACP Enabled	Active: The port automatically sends LACP protocol packets. Passive: The port does not automatically send LACP protocol packets, and responds only if it receives LACP protocol packets from the opposite device. A link having either two active LACP ports or one active port and one passive port can perform dynamic LACP trunking. A link having two passive LACP ports will not perform LACP trunking because both ports are waiting for an LACP protocol packet from the opposite device. Note: In the default switch configuration, LACP is disabled for all ports.
Trunk Group	TrkX: This port has been manually configured into a static LACP trunk. Trunk Group Same as Port Number: The port is configured for LACP, but is not a member of a port trunk.
Port Status	Up: The port has an active LACP link and is not blocked or in Standby mode. Down: The port is enabled, but an LACP link is not established. This can indicate, for example, a port that is not connected to the network or a speed mismatch between a pair of linked ports. Disabled: The port cannot carry traffic. Blocked: LACP, spanning tree has blocked the port. (The port is not in LACP Standby mode.) This may be due to a (brief) trunk negotiation or a configuration error such as differing port speeds on the same link or trying to connect the switch to more trunks than it can support. (See the table on page 12-6.) Note: Some older devices are limited to four ports in a trunk. When eight LACP-enabled ports are connected to one of these older devices, four ports connect, but the other four ports are blocked. Standby: The port is configured for dynamic LACP trunking to another device, but the maximum number of ports for the Dynamic trunk to that device has already been reached on either the switch or the other device. This port will remain in reserve, or “standby” unless LACP detects that another, active link in the trunk has become disabled, blocked, or down. In this case, LACP automatically assigns a Standby port, if available, to replace the failed port.

Status Name	Meaning
LACP Partner	Yes: LACP is enabled on both ends of the link. No: LACP is enabled on the switch, but either LACP is not enabled or the link has not been detected on the opposite device.
LACP Status	Success: LACP is enabled on the port, detects and synchronizes with a device on the other end of the link, and can move traffic across the link. Failure: LACP is enabled on a port and detects a device on the other end of the link, but is not able to synchronize with this device, and therefore not able to send LACP packets across the link. This can be caused, for example, by an intervening device on the link (such as a hub), a bad hardware connection, or if the LACP operation on the opposite device does not comply with the IEEE 802.3ad standard.

LACP Notes and Restrictions

802.1X (Port-Based Access Control) Configured on a Port. 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.

```
ProCurve(config)# aaa port-access authenticator b1
LACP has been disabled on 802.1x port(s).
ProCurve(config)#
```

The switch will not allow you to configure LACP on a port on which port access (802.1X) is enabled. For example:

```
ProCurve(config)# int b1 lacp passive
Error configuring port < port-number >: LACP and 802.1x
cannot be run together.
ProCurve(config)#
```

To restore LACP to the port, you must first remove the port's 802.1X configuration and then re-enable LACP active or passive on the port.

Port Security Configured on a Port. 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 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 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.

Changing Trunking Methods. To convert a trunk from static to dynamic, you must first eliminate the static trunk.

Static LACP Trunks. Where a port is configured for LACP (Active or Passive), but does not belong to an existing trunk group, you can add that port to a static trunk. Doing so disables dynamic LACP on that port, which means you must manually configure both ends of the trunk.

Dynamic LACP Trunks. You can configure a port for LACP-active or LACP-passive, but on a dynamic LACP trunk you cannot configure the other options that you can on static trunks. If you want to manually configure a trunk, use the **trunk** command. (Refer to “Using the CLI To Configure a Static or Dynamic Trunk Group” on page 12-15.)

VLANs and Dynamic LACP. A dynamic LACP trunk operates only in the default VLAN (unless you have enabled GVRP on the switch and use **Forbid** to prevent the ports from joining the default VLAN).

- If you want to use LACP for a trunk on a non-default VLAN and GVRP is disabled, configure the trunk as a static trunk.

Blocked Ports with Older Devices. Some older devices are limited to four ports in a trunk. When eight LACP-enabled ports are connected to one of these older devices, four ports connect, but the other four ports are blocked. The LACP status of the blocked ports is shown as “Failure”.

If one of the other ports becomes disabled, a blocked port will replace it (Port Status becomes “Up”). When the other port becomes active again, the replacement port goes back to blocked (Port Status is “Blocked”). It can take a few seconds for the switch to discover the current status of the ports.


```
ProCurve(eth-B1-B8)# show lacp

                                LACP

PORT      LACP      TRUNK      PORT      LACP      LACP
NUMB     ENABLED   GROUP      STATUS    PARTNER   STATUS
-----
B1       Active   Dyn1       Up        Yes       Success
B2       Active   Dyn1       Up        Yes       Success
B3       Active   Dyn1       Up        Yes       Success
B4       Active   Dyn1       Up        Yes       Success
B5       Active   Dyn1       Blocked  Yes       Failure
B6       Active   Dyn1       Blocked  Yes       Failure
B7       Active   B7        Down     No        Success
B8       Active   B8        Down     No        Success
```

Figure 12-11. Blocked Ports with LACP

- If there are ports that you do not want on the default VLAN, ensure that they cannot become dynamic LACP trunk members. Otherwise a traffic loop can unexpectedly occur. For example:

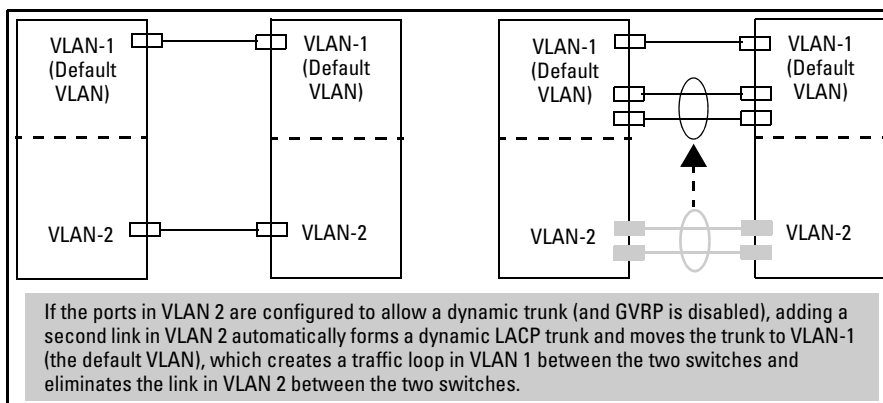


Figure 12-12. A Dynamic LACP Trunk Forming in a VLAN Can Cause a Traffic Loop

Easy control methods include either disabling LACP on the selected ports or configuring them to operate in static LACP trunks.

Spanning Tree and IGMP. If Spanning Tree and/or IGMP is enabled in the switch, a dynamic LACP trunk operates only with the default settings for these features and does not appear in the port listings for these features.

Half-Duplex and/or Different Port Speeds Not Allowed in LACP

Trunks. The ports on both sides of an LACP trunk must be configured for the same speed and for full-duplex (FDx). The 802.3ad LACP standard specifies a full-duplex (FDx) requirement for LACP trunking. (10-gigabit ports operate only at FDx.)

A port configured as LACP passive and not assigned to a port trunk can be configured to half-duplex (HDx). However, in any of the following cases, a port cannot be reconfigured to an HDx setting:

- If the port is a 10-gigabit port.
- If a port is set to LACP Active, you cannot configure it to HDx.
- If a port is already a member of a static or dynamic LACP trunk, you cannot configure it to HDx.
- If a port is already set to HDx, the switch does not allow you to configure it for a static or dynamic LACP trunk.

Dynamic/Static LACP Interoperation: A port configured for dynamic LACP can properly interoperate with a port configured for static (TrkX) LACP, but any ports configured as standby LACP links will be ignored.

Distributed Trunking

Overview

The IEEE standard 802.3ad requires that all the links in a trunk group originate from the same switch. Distributed Trunking uses a proprietary protocol that allows two or more port trunk links distributed across two switches to create a trunk group. The grouped links appear to the downstream device as if they are from a single device. This allows third party devices to interoperate with the Distributed Trunking switches (DTSs) seamlessly. Distributed trunking also provides node-level Layer 2 resiliency in a Layer 2 network if one of the switches fails.

Distributed trunking switches are connected by a special interface called the InterSwitch-Connect (ISC) port. This interface exchanges information so that the DTSs appear as a single switch to a downstream device.

The downstream device is a Distributed Trunking Device (DTD). Only servers are supported as DTDs. The DTD (server) forms a trunk with the DTSs. The connecting links are DT links and the ports are DT ports. A Distributed Trunk can span a maximum of two switches.

Distributed trunks can be grouped together by configuring two individual dt-lacp trunks with the same trunk group name in each switch. The DT ports are grouped dynamically after the configuration of distributed trunking. Figure 12-13 shows a basic distributed trunking configuration.

Note

Before you configure the switch, it is recommended that you review the “Distributed Trunking Restrictions” on page 12-35 for a complete list of operating notes and restrictions.

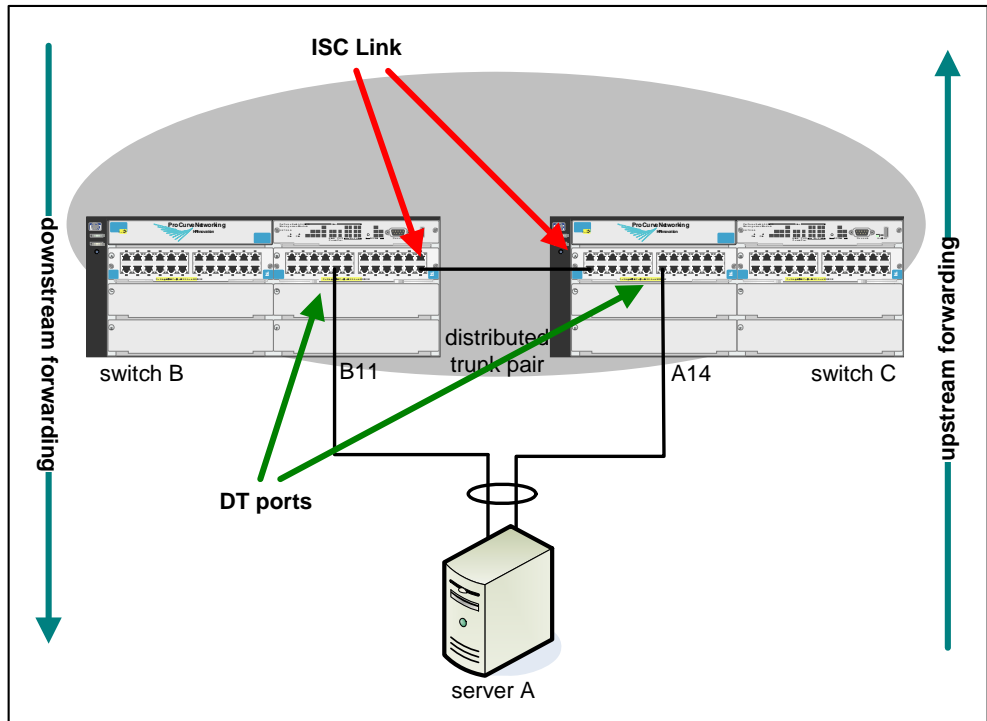


Figure 12-13. Example of Distributed Trunking Configuration

In figure 12-14, three different distributed trunks with three different servers have one common ISC link. Each trunk only spans two distributed trunking switches. The distributed trunking switches are connected at the ISC ports so they can exchange information that allows them to appear as one device to the server.

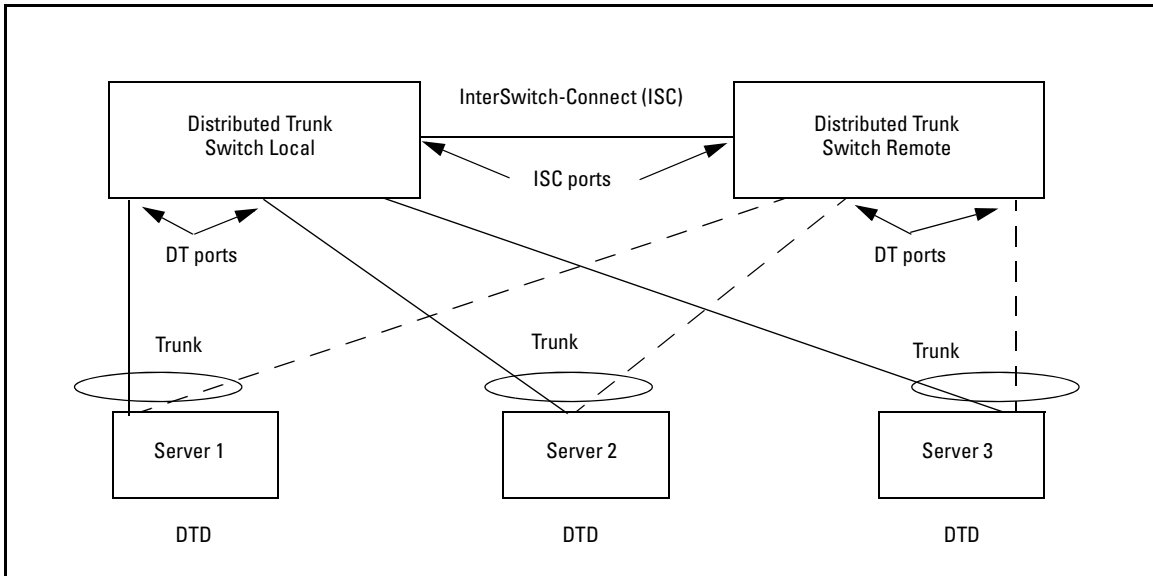


Figure 12-14. Example of Distributed Trunking

Distributed Trunking Interconnect Protocol (DTIP)

Distributed trunking uses the Distributed Trunking Interconnect Protocol (DTIP), which sends and receives proprietary protocol frames between two DT switches. DT ports are treated as standard LACP trunk ports. LACPDU s are advertised through DT ports.

DTIP interacts with LACP to synchronize the information about grouped links and the link status of the other distributed trunking switch.

Configuring Distributed Trunking

ISC Port Configuration

You must configure the ISC ports before you can configure the trunks for distributed trunking. To configure an ISC port, enter this command:

Syntax: switch-interconnect <port-num | trk1...trkN>
no switch-interconnect

Configures an InterSwitch-Connection (ISC) port. The <port-num | trk1...trkN> variable is the interconnect interface that connects two distributed trunking switches. It can be a physical port, manual LACP trunk, or manual non-protocol trunk. You can override an ISC configuration by configuring the command with a different value.

*The **no** form of the command removes the ISC interface configuration.*

Note: *A port that is already part of a trunk can't be configured as an ISC interface.*

Distributed Trunking Port Configuration

Distributed trunking ports must be configured manually.

To configure distributed trunking on the switch, enter this command:

Syntax: trunk <port-list> <trk1 | trk2 ... trkN> [trunk | lacp | dt-lacp]
no trunk <port-list>

*Configures distributed trunking on a switch. Use the **dt-lacp** option. The trunk groups must be identical in both switches, for example, if Switch Local is configured with trk1 and uses the **dt-lacp** option, Switch Remote must be configured with trk1 and use the **dt-lacp** option also in order to form a Distributed Trunk.*

*The **no** form of the command removes the distributed trunking configuration on the switch.*

The example in figure 12-15 shows an ISC port being configured for the Local switch and the Remote switch.

```
ProCurve Switch Local(config)# switch-interconnect a7
ProCurve Switch Remote(config)# switch-interconnect a8

ProCurve Switch Local(config)# trunk a9-a10 trk10 dt-lacp
ProCurve Switch Remote(config)# trunk a5-a6 trk10 dt-lacp
```

Figure 12-15. Example of Configuring Distributed Trunking

Displaying Distributed Trunking Information

To display information about the distributed trunks, enter the **show lacp distributed** command.

Syntax: show lacp [distributed]

Displays information about distributed trunks and LACP status.

```
ProCurve Switch Local(config)# show lacp distributed

                                DISTRIBUTED LACP

Local Port Status:

  PORT  LACP    TRUNK  PORT  LACP    LACP
  NUMB  ENABLED GROUP  STATUS PARTNER STATUS
  -----
  A9    Active  Trk10  Up    Yes     Success
  A10   Active  Trk10  Up    Yes     Success

Remote Port Status:

  PORT  LACP    TRUNK  PORT  LACP    LACP
  NUMB  ENABLED GROUP  STATUS PARTNER STATUS
  -----
  A5    Active  Trk10  Up    Yes     Success
  A6    Active  Trk10  Up    Yes     Success
```

Figure 12-16. Example of the Output for the show lacp distributed Command

Maximum DT Trunks and Links Supported

Table 12-1 shows the maximum number of DT trunks and DT links that are supported.

Table 12-1. Maximum DT Trunks and Links

Description	Max Number
Maximum number of groups (DT trunks) in a DT switch (that is, maximum number of servers supported)	60
Maximum number of switches that can be aggregated	2
Maximum number of physical links that can be aggregated in a single switch from a server (that is, maximum number of ports that can be in a trunk connected to a single switch)	4

From the server perspective, this means that there could be a maximum total of 60 servers connected to two DT switches. Each server can have up to four physical links aggregated in a single switch, meaning that a single server could have a maximum of eight links (that is, four on each DT switch) in a DT trunk.

Forwarding Traffic with Distributed Trunking and Spanning Tree

Refer to figure 12-17 for the following discussion about forwarding traffic when spanning tree is enabled. In this example, it is assumed that traffic is sent through switch B from a host to a server and from the server back to the host. STP can block any one of the upstream links; in this example STP has blocked all the links except the I1 link connected to DT1.

Forwarding Unicast Traffic Upstream

The server uses load balancing when it sends traffic to a DT switch. The load balancing is usually based on a SA/DA (source address/destination address) combination. See “Outbound Traffic Distribution Across Trunked Links” on page 12-37 for more information.

If unicast traffic is received on switch DT1, and the switch knows the destination MAC address of any of the devices on the upstream links, such as on link I1, it forwards the unicast frames upstream using link I1. If switch DT2 receives the unicast traffic from the server, the destination MAC address was already learned on the DT1 link, so the traffic is forwarded to the DT1 switch, which then forwards the traffic upstream on link I1.

Unicast frames are only forwarded by one of the DT switches unless the MAC address is reachable only through the other DT switch, for example, a host on DT2 sends or receives frames directly through the DT2 switch.

Forwarding Broadcast, Multicast, and Unknown Traffic Upstream

When the DT1 switch receives broadcast traffic, multicast traffic, or traffic with an unknown destination from a server, it broadcasts the traffic across all the links, including the DT1 link to the DT2 switch. The DT2 switch won't send this traffic through the same trunk groups, so the server doesn't receive the same traffic that it sent out. The DT2 switch only forwards this traffic to switches or hosts that are not on its DT links.

Forwarding Unicast Traffic Downstream (to the Server)

In normal trunking, the load balancing is based on the SA/DA MAC address pair across the links in the trunk (See "Outbound Traffic Distribution Across Trunked Links" on page 12-37 for more information on SA/DA load balancing). The last 5 bits of the SA/DA MAC address pair are used in a calculation to determine the load-balancing across the links. In distributed trunking, this calculation is extended across the DT switches and its links.

Downstream unicast traffic received by the DT switches is load balanced across both switches using the DT trunks, based on the SA/DA MAC address pair. For example, if the DT1 switch and the DT2 switch have the same number of DT links operational in a trunk group (1 link on DT1 and 1 link on DT2, 2 links on DT1 and 2 links on DT2, etc.), half of the traffic is forwarded by the DT1 switch and half of the traffic is forwarded by the DT2 switch. This is not the same as forwarding half of the *bandwidth* traffic. If DT1 has one link and DT2 has two links, the SA/DA traffic flow is load balanced with 33% of the traffic flowing through DT1 and 66% of the traffic flowing through DT2.

Forwarding Broadcast, Multicast, and Unknown Traffic Downstream (to the Server)

Downstream broadcast traffic, multicast traffic, and traffic with an unknown destination is sent the same way that unicast traffic is sent, that is, the DT1 switch and DT2 switch forward the traffic proportionately based on the number of DT links in each switch in a trunk group.

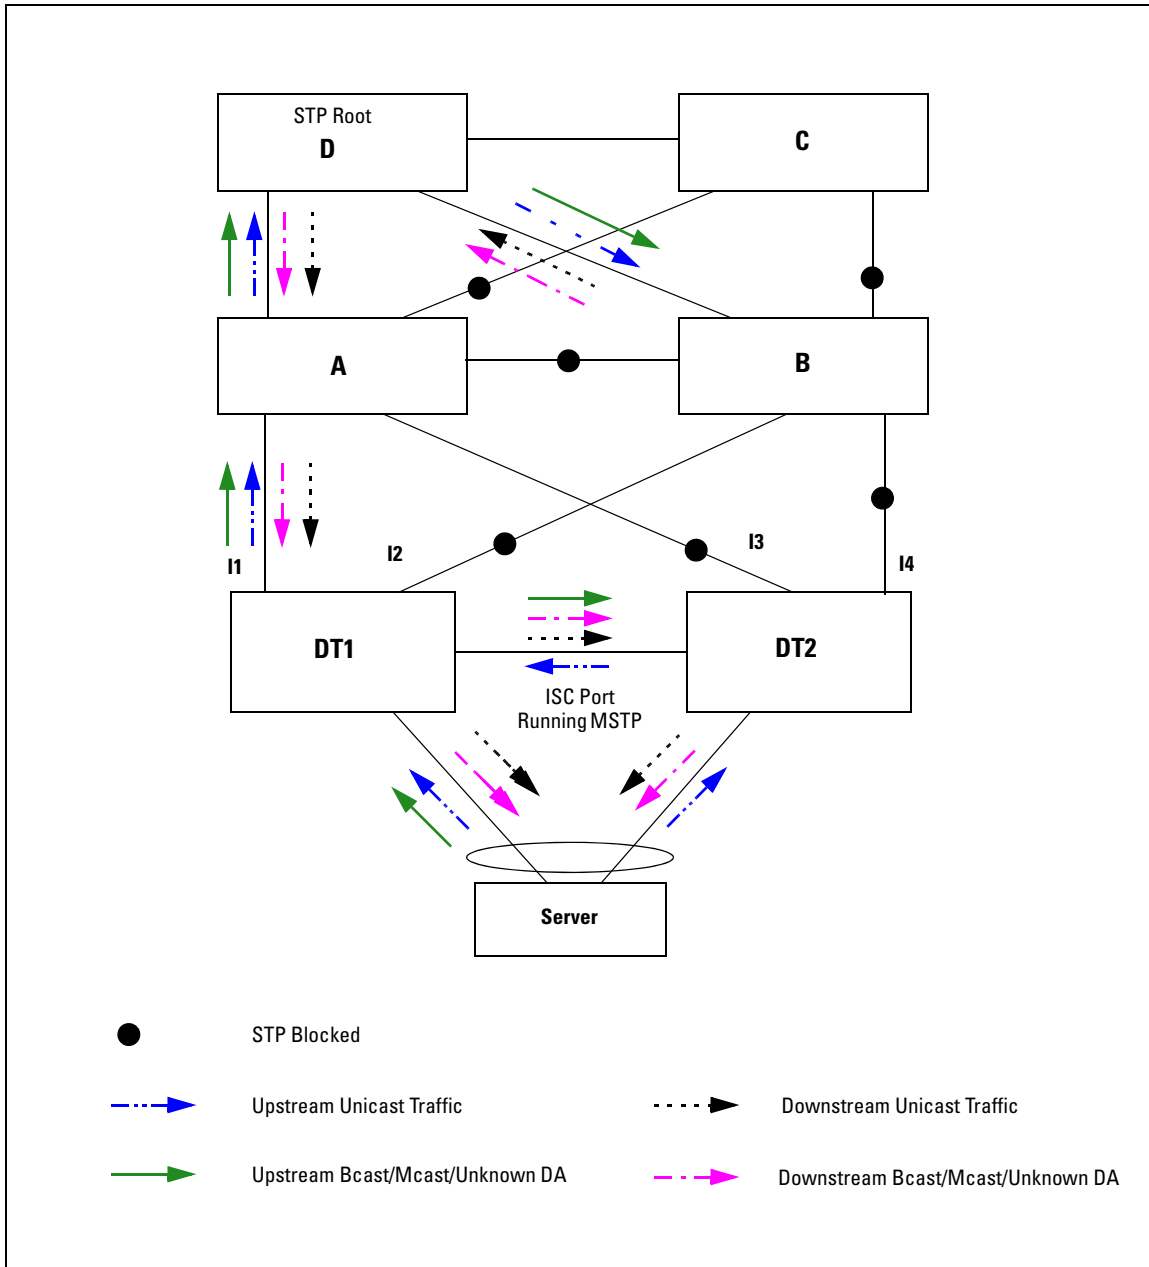


Figure 12-17. Example of Distributed Trunking with STP Forwarding Unicast, Broadcast and Multicast Traffic

Distributed Trunking Restrictions

There are several restrictions with distributed trunking.

- The port trunk links should be configured manually (manual LACP). Dynamic linking across switches is not supported.
- Only servers are supported as Distributed Trunking Devices (DTDs).
- A distributed trunk can span a maximum of two switches.
- A maximum total of 60 servers can be connected to two DT switches. Each server can have up to four physical links aggregated in a single switch, meaning that there can be a maximum of eight ports (four aggregated links for each DT switch) included in a DT trunk.
- Only one ISC link is supported per switch with a maximum of 60 DT trunks supported on the switch. The ISC link can be configured as a manual LACP trunk, non-protocol trunk, or as an individual link. Dynamic LACP trunks are not supported as ISCs.
- An ISC port becomes a member of all VLANs that are configured on the switch. When a new VLAN is configured, the ISC ports become members of that VLAN.
- Port trunk links can be done only on a maximum of two switches that are connected to a specific server.
- Any VLAN that is in a distributed trunk must be configured on both switches. By default, the distributed trunk belongs to the default VLAN.
- There can be eight links in a distributed trunk grouped across two switches, with a limit of four links per distributed trunking switch.
- The limit of 60 manual trunks per switch includes distributed trunking manual trunks as well.
- IP routing and distributed trunking are mutually exclusive. Routing restrictions with distributed trunking are switch-wide and do not apply to the DT ports only.
- Meshing and DT switches are mutually exclusive.
- ARP protection is not supported on the distributed trunks.
- STP is disabled on DT ports.
- QinQ in mixed VLAN mode and distributed trunking are mutually exclusive.
- SVLANs in mixed mode are not supported on DT or ISC links.
- DHCP snooping and IGMP snooping are not supported on DT links.

Trunk Group Operation Using the “Trunk” Option

This method creates a trunk group that operates independently of specific trunking protocols and does not use a protocol exchange with the device on the other end of the trunk. With this choice, the switch simply uses the SA/DA method of distributing outbound traffic across the trunked ports without regard for how that traffic is handled by the device at the other end of the trunked links. Similarly, the switch handles incoming traffic from the trunked links as if it were from a trunked source.

When a trunk group is configured with the **trunk** option, the switch automatically sets the trunk to a priority of “4” for spanning-tree operation (even if spanning-tree is currently disabled. This appears in the running-config file as `spanning-tree Trkn priority 4`. Executing **write memory** after configuring the trunk places the same entry in the startup-config file.

Use the Trunk option to establish a trunk group between a switch covered in this guide and another device, where the other device’s trunking operation fails to operate properly with LACP trunking configured on the switches.

How the Switch Lists Trunk Data

Static Trunk Group: Appears in the menu interface and the output from the CLI **show trunk** and **show interfaces** commands.

Dynamic LACP Trunk Group: Appears in the output from the CLI **show lacp** command.

Interface Option	Dynamic LACP Trunk Group	Static LACP Trunk Group	Static Non-Protocol
Menu Interface	No	Yes	Yes
CLI show trunk	No	Yes	Yes
CLI show interfaces	No	Yes	Yes
CLI show lacp	Yes	Yes	No
CLI show spanning-tree	No	Yes	Yes
CLI show igmp	No	Yes	Yes
CLI show config	No	Yes	Yes

Outbound Traffic Distribution Across Trunked Links

The two trunk group options (LACP and Trunk) use source-destination address pairs (SA/DA) for distributing outbound traffic over trunked links. SA/DA (source address/destination address) causes the switch to distribute outbound traffic to the links within the trunk group on the basis of source/destination address pairs. That is, the switch sends traffic from the same source address to the same destination address through the same trunked link, and may also send traffic from the same source address to a different destination address through the same link or a different link, depending on the mapping of path assignments among the links in the trunk. Likewise, the switch distributes traffic for the same destination address but from different source addresses through links depending on the path assignment.

The load-balancing is done on a per communication basis. Otherwise, traffic is transmitted across the same path as shown in figure 12-18. That is, if Client A attached to Switch 1 sends five packets of data to Server A attached to Switch 2, the same link is used to send all five packets. The SA/DA address pair for the traffic is the same. The packets are not evenly distributed across any other existing links between the two switches; they all take the same path.

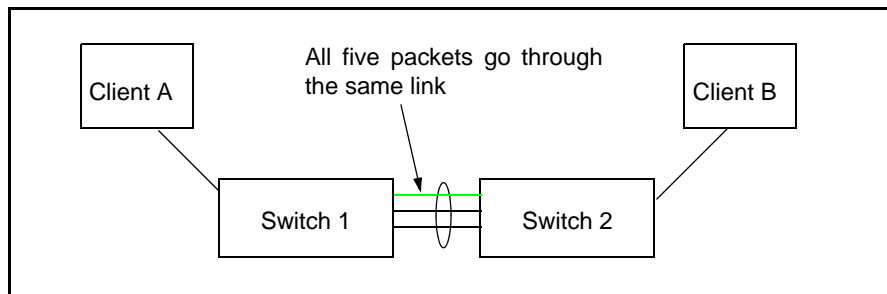


Figure 12-18. Example of Single Path Traffic through a Trunk

The actual distribution of the traffic through a trunk depends on a calculation using bits from the Source Address and Destination address. When an IP address is available, the calculation includes the last five bits of the IP source address and IP destination address, otherwise the MAC addresses are used. The result of that process undergoes a mapping that determines which link the traffic goes through. If you have only two ports in a trunk, it is possible that all the traffic will be sent through one port even if the SA/DA pairs are different. The more ports you have in the trunk, the more likely it is that the traffic will be distributed among the links.

When a new port is added to the trunk, the switch begins sending traffic, either new traffic or existing traffic, through the new link. As links are added or deleted, the switch redistributes traffic across the trunk group. For example, in figure 12-19 showing a three-port trunk, traffic could be assigned as shown in table 12-1.




Figure 12-19. Example of Port-Trunked Network

Table 12-1. Example of Link Assignments in a Trunk Group (SA/DA Distribution)

Source:	Destination:	Link:
Node A	Node W	1
Node B	Node X	2
Node C	Node Y	3
Node D	Node Z	1
Node A	Node Y	2
Node B	Node W	3

Because the amount of traffic coming from or going to various nodes in a network can vary widely, it is possible for one link in a trunk group to be fully utilized while other links in the same trunk have unused bandwidth capacity even if the assignments were evenly distributed across the links in a trunk.

Port Trunking
Outbound Traffic Distribution Across Trunked Links

Port Traffic Controls

Contents

Overview	13-3
Rate-Limiting	13-4
All Traffic Rate-Limiting	13-4
Configuring Rate-Limiting	13-5
Displaying the Current Rate-Limit Configuration	13-6
Operating Notes for Rate-Limiting	13-8
ICMP Rate-Limiting	13-10
Terminology	13-11
Guidelines for Configuring ICMP Rate-Limiting	13-12
Configuring ICMP Rate-Limiting	13-13
Using Both ICMP Rate-Limiting and All-Traffic Rate-Limiting on the Same Interface	13-14
Displaying the Current ICMP Rate-Limit Configuration	13-14
Operating Notes for ICMP Rate-Limiting	13-15
ICMP Rate-Limiting Trap and Event Log Messages	13-17
Configuring Inbound Rate-Limiting for Broadcast and Multicast Traffic	13-19
Operating Notes	13-21
Guaranteed Minimum Bandwidth (GMB)	13-22
Introduction	13-22
Terminology	13-22
GMB Operation	13-22
Impacts of QoS Queue Configuration on GMB Operation	13-24
Configuring Guaranteed Minimum Bandwidth for Outbound Traffic	13-25
Displaying the Current Guaranteed Minimum Bandwidth Configuration	13-28
GMB Operating Notes	13-29

Jumbo Frames	13-30
Terminology	13-30
Operating Rules	13-31
Configuring Jumbo Frame Operation	13-32
Overview	13-32
Viewing the Current Jumbo Configuration	13-33
Enabling or Disabling Jumbo Traffic on a VLAN	13-35
Configuring a Maximum Frame Size	13-35
Configuring IP MTU	13-36
SNMP Implementation	13-36
Displaying the Maximum Frame Size	13-37
Operating Notes for Maximum Frame Size	13-37
Operating Notes for Jumbo Traffic-Handling	13-37
Troubleshooting	13-40

Overview

Feature	Default	Menu	CLI	Web
Rate-Limiting	None	n/a	13-4	n/a
Guaranteed Minimum Bandwidth	Per Queue (1-8 order): 2%-3%-30%-10%-10%- 10%-15%-20%	n/a	13-22	n/a
Jumbo Packets	Disabled	n/a	13-30	n/a

This chapter includes:

- **Rate-Limiting:** Enables a port to limit the amount of bandwidth a user or device may utilize for traffic on the switch.

Note

In earlier releases, all traffic rate-limiting applied to inbound traffic only, and was specified as a percentage of total bandwidth. Beginning with software release K.12.*xxx* or later, it is also possible to configure outbound rate-limiting for all traffic on a port, and specify bandwidth usage in terms of kilobits per second (kbps).

- **Guaranteed Minimum Bandwidth (GMB):** Provides a method for ensuring that each of a port's outbound queues has a specified minimum consideration for sending traffic out on the link to another device.
- **Jumbo Frames:** Enables ports operating at a minimum of 10 Mbps on the ProCurve 3500 switches and 1 Gbps on the other switches covered in this guide to accept inbound frames of up to 9220 bytes when configured for jumbo traffic.

Rate-Limiting

Feature	Default	Menu	CLI	Web
rate-limit all	none	n/a	page 13-5	n/a
show rate-limit all	n/a	n/a	page 13-6	n/a
rate-limit icmp	none	n/a	page 13-13	n/a
show rate-limit icmp	n/a	n/a	page 13-14	n/a

All Traffic Rate-Limiting

Rate-limiting for all traffic operates on a per-port basis to allow only the specified bandwidth to be used for inbound or outbound traffic. When traffic exceeds the configured limit, it is dropped. This effectively sets a usage level on a given port, and is a tool for enforcing maximum service level commitments granted to network users. This feature operates on a per-port level and is not configurable on port trunks. Note that rate-limiting is designed to be applied at the network edge to limit traffic from non-critical users or to enforce service agreements such as those offered by Internet Service Providers (ISPs) to provide only the bandwidth for which a customer has paid.

Note

Rate-limiting also can be applied by a RADIUS server during an authentication client session. For further details, refer to the chapter titled “*RADIUS Authentication and Accounting*” in the *Access Security Guide* for your switch.

Caution

Rate-limiting is intended for use on edge ports in a network. It is not recommended for use on links to other switches, routers, or servers within a network, or for use in the network core. Doing so can interfere with applications the network requires to function properly.

Note

The switches covered in this guide also support ICMP rate-limiting to mitigate the effects of certain ICMP-based attacks. For more information, refer to “ICMP Rate-Limiting” on page 13-10.

Configuring Rate-Limiting

Note

The mode using bits per second (bps) in releases before K.12.XX has been replaced by the kilobits per second (kbps) mode. Switches that have configurations with bps values will be automatically converted when you update your software to the new version. However, an older config file with bps values must be updated manually to kbps values or it will not load successfully onto a switch running later versions of the software (K.12.XX or greater).

The **rate-limit all** command controls the rate of traffic sent or received on a port by setting a limit on the bandwidth available. It includes options for:

- Rate-limiting on either inbound or outbound traffic.
- Specifying the traffic rate as either a percentage of bandwidth, or in terms of bits per second.

Syntax: [no] int <port-list> rate-limit all < in | out > <percent <0-100> | kbps < 0-10000000>>

Configures a traffic rate limit (on non-trunked ports) on the link. The “no” form of the command disables rate-limiting on the specified ports.

*(Default: **Disabled.**)*

Options include:

- **in** or **out** — Specifies a traffic rate limit on inbound traffic passing through that port, or on outbound traffic.
- **percent** or **kbps** — Specifies the rate limit as a percentage of total available bandwidth, or in kilobits per second.

Notes:

- The **rate-limit icmp** command specifies a rate limit on inbound ICMP traffic only (see “ICMP Rate-Limiting” on page 13-9).
- Rate-limiting does not apply to trunked ports (including meshed ports).

—Continued—

- *Kbps rate-limiting is done in segments of 1% of the lowest corresponding media speed. For example, if the media speed is 100 Kbps, the value would be 1 Mbps. A 1-100 Kbps rate-limit is implemented as a limit of 100 Kbps; a limit of 100-199 Kbps is also implemented as a limit of 100 Kbps, a limit of 200-299 Kbps is implemented as a limit of 200 Kbps, and so on.*
- *Percentage limits are based on link speed. For example, if a 100 Mbps port negotiates a link at 100 Mbps and the inbound rate-limit is configured at 50%, then the traffic flow through that port is limited to no more than 50 Mbps. Similarly, if the same port negotiates a 10 Mbps link, then it allows no more than 5 Mbps of inbound traffic.*

*Configuring a rate limit of 0 (zero) on a port blocks all traffic on that port. However, if this is the desired behavior on the port, ProCurve recommends using the **< port-list > disable** command instead of configuring a rate limit of 0.*

You can configure a rate limit from either the global configuration level or from the port context level. For example, either of the following commands configures an inbound rate limit of 60% on ports A3 - A5:

```
ProCurve (config)# int a3-a5 rate-limit all in percent 60
ProCurve (eth-A3-A5)# rate-limit all in percent 60
```

Displaying the Current Rate-Limit Configuration

The **show rate-limit all** command displays the per-port rate-limit configuration in the running-config file.

Syntax: show rate-limit all [*port-list*]

*Without [**port-list**], this command lists the rate-limit configuration for all ports on the switch. With [**port-list**], this command lists the rate-limit configuration for the specified port(s). This command operates the same way in any CLI context.*

For example, if you wanted to view the rate-limiting configuration on the first six ports in the module in slot “A”:

```
ProCurve# show rate-limit all a1-a6
```

Ports A1-A4 are configured with an outbound rate limit of 200 Kbps; Port A5 is configured with an inbound rate limit of 20%. (Port A6 is not configured for rate-limiting.)

All-Traffic Rate Limit Maximum %

Port	Inbound			Outbound		
	Limit	Mode	Radius Override	Limit	Mode	Radius Override
A1	Disabled	Disabled	No-override	200	kbps	No-override
A2	Disabled	Disabled	No-override	200	kbps	No-override
A3	Disabled	Disabled	No-override	200	kbps	No-override
A4	Disabled	Disabled	No-override	200	kbps	No-override
A5	20	%	No-override	Disabled	Disabled	No-override
A6	Disabled	Disabled	No-override	Disabled	Disabled	No-override

Figure 13-1. Example of Listing the Rate-Limit Configuration

Note

To view RADIUS-assigned rate-limit information, use one of the following command options:

```
show port-access
  web-based clients < port-list > detailed
  mac-based clients < port-list > detailed
  authenticator clients < port-list > detailed
```

For more on RADIUS-assigned rate-limits, refer to the chapter titled “Configuring RADIUS Server Support for Switch Services” in the latest Management and Configuration Guide for your switch.

The **show running** command displays the currently applied setting for any interfaces in the switch configured for all traffic rate-limiting and ICMP rate limiting. The **show config** command displays this information for the configuration currently stored in the startup-config file. (Note that configuration changes performed with the CLI, but not followed by a **write mem** command do not appear in the startup-config file.)

```
ProCurve(config)# show config

Startup configuration:

; J8697A Configuration Editor; Created on release #K.14.01

hostname "ProCurve Switch 8212z1"
module 1 type J8705A
snmp-server community "public" Unrestricted
vlan 1
  name "DEFAULT_VLAN"
  untagged A1-A24
  ip address dhcp-bootp
  exit
interface A1
  rate-limit all out kbps 200
  exit
interface A2
  rate-limit all out kbps 200
  exit
interface A3
  rate-limit all out kbps 200
  exit
interface A4
  rate-limit all out kbps 200
  exit
interface A5
  rate-limit all in percent 200
  exit
interface A6
  rate-limit icmp percent 60
  rate-limit mcast in percent 60
  exit
```



Figure 13-2. Example of Rate-Limit Settings Listed in the “show config” Output

Operating Notes for Rate-Limiting

- **Rate-limiting operates on a per-port basis**, regardless of traffic priority. Rate-limiting is available on all types of ports (other than trunked ports) on the switches covered in this guide, and at all port speeds configurable for these devices.
- **Rate-limiting is not allowed on trunked ports:** Rate-limiting is not supported on ports configured in a trunk group (including mesh ports). Configuring a port for rate-limiting and then adding it to a trunk suspends

rate-limiting on the port while it is in the trunk. Attempting to configure rate-limiting on a port that already belongs to a trunk generates the following message:

```
<port-list>: Operation is not allowed for a trunked port.
```

- **Rate-limiting for inbound and outbound traffic are separate features:** The rate limits for each direction of traffic flow on the same port are configured separately—even the specified limits can be different.
- **Rate-limiting is visible as an outbound forwarding rate:** Because inbound rate-limiting is performed on packets during packet-processing, it is not shown via the inbound drop counters. Instead, this limit is verifiable as the ratio of outbound traffic from an inbound rate-limited port versus the inbound rate. For outbound rate-limiting, the rate is visible as the percentage of available outbound bandwidth (assuming that the amount of requested traffic to be forwarded is larger than the rate-limit).
- **Operation with other features:** Configuring rate-limiting on a port where other features affect port queue behavior (such as flow control) can result in the port not achieving its configured rate-limiting maximum. For example, in a situation where flow control is configured on a rate-limited port, there can be enough “back pressure” to hold high-priority inbound traffic from the upstream device or application to a rate that is lower than the configured rate limit. In this case, the inbound traffic flow does not reach the configured rate and lower priority traffic is not forwarded into the switch fabric from the rate-limited port. (This behavior is termed “head-of-line blocking” and is a well-known problem with flow-control.) In another type of situation, an outbound port can become oversubscribed by traffic received from multiple rate-limited ports. In this case, the actual rate for traffic on the rate-limited ports may be lower than configured because the total traffic load requested to the outbound port exceeds the port’s bandwidth, and thus some requested traffic may be held off on inbound.
- **Traffic filters on rate-limited ports:** Configuring a traffic filter on a port does not prevent the switch from including filtered traffic in the bandwidth-use measurement for rate-limiting when it is configured on the same port. For example, ACLs, source-port filters, protocol filters, and multicast filters are all included in bandwidth usage calculations.
- **Monitoring (Mirroring) rate-limited interfaces:** If monitoring is configured, packets dropped by rate-limiting on a monitored interface will still be forwarded to the designated monitor port. (Monitoring shows what traffic is inbound on an interface, and is not affected by “drop” or “forward” decisions.)

- **Optimum rate-limiting operation:** Optimum rate-limiting occurs with 64-byte packet sizes. Traffic with larger packet sizes can result in performance somewhat below the configured bandwidth. This is to ensure the strictest possible rate-limiting of all sizes of packets.

Note on Testing Rate-Limiting

Rate-limiting is applied to the available bandwidth on a port, and not to any specific applications running through the port. If the total bandwidth requested by all applications is less than the configured maximum rate, then no rate-limit can be applied. This situation occurs with a number of popular throughput-testing applications, as well as most regular network applications. Consider the following example that uses the minimum packet size:

The total available bandwidth on a 100 Mbps port “X” (allowing for Inter-packet Gap—IPG), with no rate-limiting restrictions, is:

$$(((100,000,000 \text{ bits}) / 8) / 84) \times 64 = 9,523,809 \text{ bytes per second}$$

where:

- The divisor (84) includes the 12-byte IPG, 8-byte preamble, and 64-bytes of data required to transfer a 64-byte packet on a 100 Mbps link.
- Calculated “bytes-per-second” includes packet headers and data. This value is the maximum “bytes-per-second” that 100 Mbps can support for minimum-sized packets.

Suppose port “X” is configured with a rate limit of 50% (4,761,904 bytes). If a throughput-testing application is the only application using the port, and transmits 1 Mbyte of data through the port, it uses only 10.5% of the port’s available bandwidth, and the rate-limit of 50% has no effect. This is because the maximum rate permitted (50%) exceeds the test application’s bandwidth usage (126,642-164,062 bytes, depending upon packet size, which is only 1.3-1.7% of the available total). Before rate-limiting can occur, the test application’s bandwidth usage must exceed 50% of the port’s total available bandwidth. That is, to test the rate-limit setting, the following must be true:

$$\text{bandwidth usage} > (0.50 \times 9,523,809)$$

ICMP Rate-Limiting

In IP networks, ICMP (Internet Control Message Protocol) messages are generated in response to either inquiries or requests from routing and diagnostic functions. These messages are directed to the applications originating the inquiries. In unusual situations, if the messages are generated rapidly with the intent of overloading network circuits, they can threaten network availability. This problem is visible in denial-of-service (DoS) attacks or other malicious behaviors where a worm or virus overloads the network with ICMP

messages to an extent where no other traffic can get through. (ICMP messages themselves can also be misused as virus carriers). Such malicious misuses of ICMP can include a high number of ping packets that mimic a valid source IP address and an invalid destination IP address (spoofed pings), and a high number of response messages (such as Destination Unreachable error messages) generated by the network.

ICMP rate-limiting provides a method for limiting the amount of bandwidth that may be utilized for inbound ICMP traffic on a switch port or trunk. This feature allows users to restrict ICMP traffic to percentage levels that permit necessary ICMP functions, but throttle additional traffic that may be due to worms or viruses (reducing their spread and effect). In addition, ICMP rate-limiting preserves inbound port bandwidth for non-ICMP traffic.

Caution

The ICMP protocol is necessary for routing, diagnostic, and error responses in an IP network. ICMP rate-limiting is primarily used for throttling worm or virus-like behavior, and should normally be configured to allow one to five per cent of available inbound bandwidth (at 10 Mbps or 100 Mbps speeds) or 100 - 10,000 kbps (1Gbps or 10 Gbps speeds) to be used for ICMP traffic. ***This feature should not be used to remove all ICMP traffic from a network.***

Note

ICMP rate-limiting does not throttle non-ICMP traffic. In cases where you want to throttle both ICMP traffic and all other inbound traffic on a given interface, you can separately configure both ICMP rate-limiting and all-traffic rate-limiting.

Beginning with software release K.12.*xx* or later, the all-traffic rate-limiting command (**rate-limit all**) and the ICMP rate-limiting command (**rate-limit icmp**) operate differently:

- All traffic rate-limiting applies to both inbound and outbound traffic, and can be specified either in terms of a percentage of total bandwidth or in terms of bits per second;
 - ICMP rate-limiting applies only to inbound traffic, and can only be specified as a percentage of total bandwidth.
-

Terminology

All-Traffic Rate-Limiting: Applies a rate-limit to all traffic (including ICMP traffic) on an interface. For details, see “Rate-Limiting” on page 13-4.

ICMP Rate-Limiting: Applies a rate-limit to all *inbound* ICMP traffic received on an interface, but does not limit other types of inbound traffic.

Spoofed Ping: An ICMP echo request packet intentionally generated with a valid source IP address and an invalid destination IP address. Spoofed pings are often created with the intent to oversubscribe network resources with traffic having invalid destinations.

Guidelines for Configuring ICMP Rate-Limiting

Apply ICMP rate-limiting on all connected interfaces on the switch to effectively throttle excessive ICMP messaging from any source. Figure 13-3 shows an example of how to configure this for a small to mid-sized campus though similar rate-limit thresholds are applicable to other network environments. On edge interfaces, where ICMP traffic should be minimal, a threshold of 1% of available bandwidth should be sufficient for most applications. On core interfaces, such as switch-to-switch and switch-to-router, a maximum threshold of 5% should be sufficient for normal ICMP traffic. (“Normal” ICMP traffic levels should be the maximums that occur when the network is rebooting.)

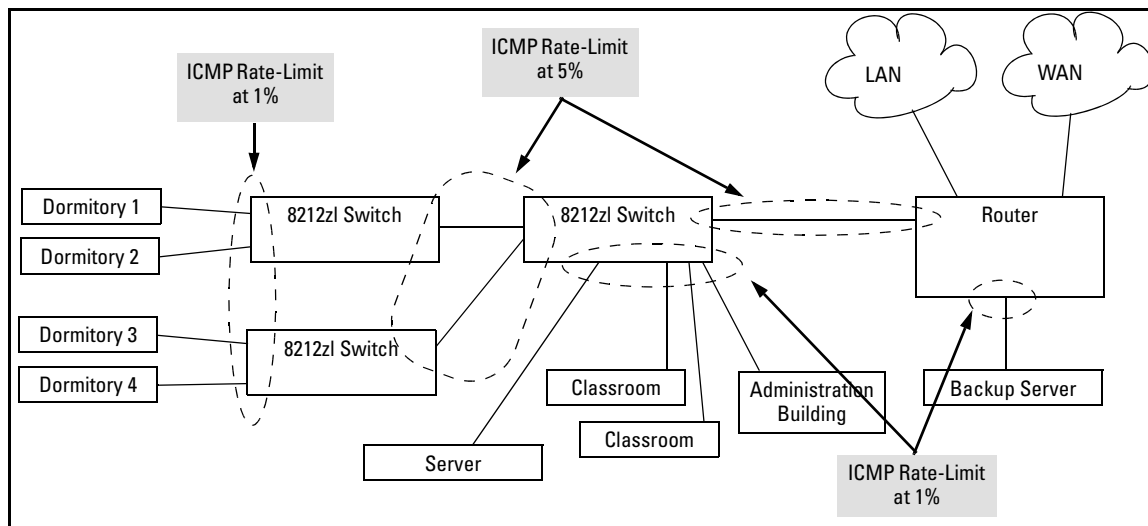


Figure 13-3. Example of ICMP Rate-Limiting

Configuring ICMP Rate-Limiting

The **rate-limit icmp** command controls inbound usage of a port by setting a limit on the bandwidth available for inbound ICMP traffic.

Syntax: [no] int < port-list > rate-limit icmp <percent < 0-100 > | kbps <0-10000000>>

*Configures inbound ICMP traffic rate limiting. You can configure a rate limit from either the global configuration level (as shown above) or from the interface context level. The **no** form of the command disables ICMP rate-limiting on the specified interface(s). (Default: **Disabled**.)*

percent <1-100>: Values in this range allow ICMP traffic as a percentage of the bandwidth available on the interface.

kbps <0-10000000>: Specifies the rate at which to forward traffic in kilobits per second.

0: This value causes an interface to drop all incoming ICMP traffic, and is not recommended. Refer to the Caution on page 13-11.

Note: ICMP rate-limiting is not supported on meshed ports. (Rate-limiting can reduce the efficiency of paths through a mesh domain).

For example, either of the following commands configures an inbound rate limit of 1% on ports A3 - A5, which are used as network edge ports:

```
ProCurve(config)# int a3-a5 rate-limit icmp 1
ProCurve (eth-A3-A5)# rate-limit icmp 1
```

Note

When using kbps-mode ICMP rate-limiting, the rate-limiting only operates on the IP payload part of the ICMP packet (as required by metering RFC 2698). This means that effective metering is at a rate greater than the configured rate, with the disparity increasing as the packet size decreases (the packet to payload ratio is higher).

Also, in kbps mode, metering accuracy is limited at low values, for example, less than 45 Kbps. This is to allow metering to function well at higher media speeds such as 10 Gbps.

Using Both ICMP Rate-Limiting and All-Traffic Rate-Limiting on the Same Interface

ICMP and all-traffic rate-limiting can be configured on the same interface. All-traffic rate-limiting applies to all inbound or outbound traffic (including ICMP traffic), while ICMP rate-limiting applies only to inbound ICMP traffic.

Note that if the all-traffic load on an interface meets or exceeds the currently configured all-traffic inbound rate-limit while the ICMP traffic rate-limit on the same interface has not been reached, then all excess traffic will be dropped, including any inbound ICMP traffic above the all-traffic limit (regardless of whether the ICMP rate-limit has been reached). Suppose, for example:

- The all-traffic inbound rate-limit on port “X” is configured at 55% of the port’s bandwidth.
- The ICMP traffic rate-limit on port “X” is configured at 2% of the port’s bandwidth.

If at a given moment:

- Inbound ICMP traffic on port “X” is using 1% of the port’s bandwidth, and
- Inbound traffic of all types on port “X” demands 61% of the ports’s bandwidth,

then all inbound traffic above 55% of the port’s bandwidth, including any additional ICMP traffic, will be dropped as long as all inbound traffic combined on the port demands 55% or more of the port’s bandwidth.

Displaying the Current ICMP Rate-Limit Configuration

The **show rate-limit icmp** command displays the per-interface ICMP rate-limit configuration in the running-config file.

Syntax: show rate-limit icmp [*port-list*]

Without [port-list], this command lists the ICMP rate-limit configuration for all ports on the switch. With [port-list], this command lists the rate-limit configuration for the specified interface(s). This command operates the same way in any CLI context.

For example, if you wanted to view the rate-limiting configuration on the first six ports in the module in slot “B”:

```
ProCurve(config)# show rate-limit icmp b1-b6

Inbound ICMP Rate Limit Maximum Percentage

Port | Mode      Rate
-----+-----
B1   | Disabled  Disabled
B2   | kbps      100
B3   | %         5
B4   | %         1
B5   | %         1
B6   | Disabled  Disabled
```

Figure 13-4. Example of Listing the Rate-Limit Configuration

The **show running** command displays the currently applied setting for any interfaces in the switch configured for all traffic rate-limiting and ICMP rate-limiting. The **show config** command displays this information for the configuration currently stored in the startup-config file. (Note that configuration changes performed with the CLI, but not followed by a **write mem** command do not appear in the startup-config file.)

Operating Notes for ICMP Rate-Limiting

ICMP rate-limiting operates on an interface (per-port) basis to allow, on average, the highest expected amount of legitimate, inbound ICMP traffic.

- **Interface support:** ICMP rate-limiting is available on all types of ports (other than trunk ports or mesh ports), and at all port speeds configurable for the switch.
- **Rate-limiting is not permitted on mesh ports:** Either type of rate-limiting (all traffic or ICMP) can reduce the efficiency of paths through a mesh domain.
- **Rate-limiting is not supported on port trunks:** Neither all-traffic nor ICMP rate-limiting are supported on ports configured in a trunk group.
- **ICMP percentage-based rate-limits are calculated as a percentage of the negotiated link speed:** For example, if a 100 Mbps port negotiates a link to another switch at 100 Mbps and is ICMP rate-limit configured at 5%, then the inbound ICMP traffic flow through that port is limited to 5 Mbps. Similarly, if the same port negotiates a 10 Mbps link, then it allows

0.5 Mbps of inbound traffic. If an interface experiences an inbound flow of ICMP traffic in excess of its configured limit, the switch generates a log message and an SNMP trap (if an SNMP trap receiver is configured).

- **ICMP rate-limiting is port-based:** ICMP rate-limiting reflects the available percentage of an interface's entire inbound bandwidth. The rate of inbound flow for traffic of a given priority and the rate of flow from an ICMP rate-limited interface to a particular queue of an outbound interface are not measures of the actual ICMP rate limit enforced on an interface.
- **Below-maximum rates:** ICMP rate-limiting operates on a per-interface basis, regardless of traffic priority. Configuring ICMP rate-limiting on an interface where other features affect inbound port queue behavior (such as flow control) can result in the interface not achieving its configured ICMP rate-limiting maximum. For example, in some situations with flow control configured on an ICMP rate-limited interface, there can be enough "back pressure" to hold high-priority inbound traffic from the upstream device or application to a rate that does not allow bandwidth for lower-priority ICMP traffic. In this case, the inbound traffic flow may not permit the forwarding of ICMP traffic into the switch fabric from the rate-limited interface. (This behavior is termed "head-of-line blocking" and is a well-known problem with flow-control.) In cases where both types of rate-limiting (**rate-limit all** and **rate-limit icmp**) are configured on the same interface, this situation is more likely to occur. In another type of situation, an outbound interface can become oversubscribed by traffic received from multiple ICMP rate-limited interfaces. In this case, the actual rate for traffic on the rate-limited interfaces may be lower than configured because the total traffic load requested to the outbound interface exceeds the interface's bandwidth, and thus some requested traffic may be held off on inbound.
- **Monitoring (Mirroring) ICMP rate-limited interfaces:** If monitoring is configured, packets dropped by ICMP rate-limiting on a monitored interface will still be forwarded to the designated monitor port. (Monitoring shows what traffic is inbound on an interface, and is not affected by "drop" or "forward" decisions.)
- **Optimum rate-limiting operation:** Optimum rate-limiting occurs with 64-byte packet sizes. Traffic with larger packet sizes can result in performance somewhat below the configured inbound bandwidth. This is to ensure the strictest possible rate-limiting of all sizes of packets.
- **Outbound Traffic Flow:** Configuring ICMP rate-limiting on an interface does not control the rate of outbound traffic flow on the interface.

**Note on Testing
ICMP Rate-Limiting**

ICMP rate-limiting is applied to the available bandwidth on an interface. If the total bandwidth requested by all ICMP traffic is less than the available, configured maximum rate, then no ICMP rate-limit can be applied. That is, an interface must be receiving more inbound ICMP traffic than the configured bandwidth limit allows. If the interface is configured with both **rate-limit all** and **rate-limit icmp**, then the ICMP limit can be met or exceeded only if the rate limit for all types of inbound traffic has not already been met or exceeded. Also, to test the ICMP limit it is necessary to generate ICMP traffic that exceeds the configured ICMP rate limit. Using the recommended settings—1% for edge interfaces and 5% maximum for core interfaces—it is easy to generate sufficient traffic. However, if you are testing with higher maximums, it is necessary to ensure that the ICMP traffic volume exceeds the configured maximum.

Note also that testing ICMP rate-limiting where inbound ICMP traffic on a given interface has destinations on multiple outbound interfaces, the test results must be based on the received outbound ICMP traffic.

ICMP rate-limiting is not reflected in counters monitoring inbound traffic because inbound packets are counted before the ICMP rate-limiting drop action occurs.

ICMP Rate-Limiting Trap and Event Log Messages

If the switch detects a volume of inbound ICMP traffic on a port that exceeds the ICMP rate-limit configured for that port, it generates one SNMP trap and one informational Event Log message to notify the system operator of the condition. (The trap and Event Log message are sent within two minutes of when the event occurred on the port.)

For example:

```
I 06/30/05 11:15:42 RateLim: ICMP traffic exceeded  
configured limit on port A1
```

These trap and Event Log messages provide an advisory that inbound ICMP traffic on a given interface has exceeded the configured maximum. The additional ICMP traffic is dropped, but the excess condition may indicate an infected host (or other traffic threat or network problem) on that interface. The system operator should investigate the attached devices or network conditions further.

The switch does not send more traps or Event Log messages for excess ICMP traffic on the affected port until the system operator resets the port's ICMP trap function. The reset can be done through SNMP from a network management station or through the CLI with the following **setmib** command.

Syntax: `setmib hpicmpRatelimitPortAlarmflag.< internal-port-#> -i 1`

On a port configured with ICMP rate-limiting, this command resets the ICMP trap function, which allows the switch to generate a new SNMP trap and an Event Log message if ICMP traffic in excess of the configured limit is detected on the port.

For example, an operator noticing an ICMP rate-limiting trap or Event Log message originating with port A1 on a switch would use the following **setmib** command to reset the port to send a new message if the condition occurs again.

```
ProCurve(config)# setmib hpicmpratelimitportalarm-  
flag.1 -i 1
```

Determining the Switch Port Number Used in ICMP Port Reset

Commands: To enable excess ICMP traffic notification traps and Event Log messages, use the **setmib** command described on page 13-17. The port number included in the command corresponds to the internal number the switch maintains for the designated port, and not the port's external (slot/number) identity.

To match the port's external slot/number to the internal port number, use the **walkmib ifDescr** command, as shown in the following figure:

```

ProCurve# walkmib ifDescr
┌ ifDescr.1 = A1 ───┐
│ ifDescr.2 = A2   │
│ ifDescr.3 = A3   │
│ .               │
│ .               │
│ ifDescr.23 = A23 │
│ ifDescr.24 = A24 │
│ ifDescr.27 = B1   │
│ ifDescr.28 = B2   │
│ ifDescr.29 = B3   │
│ .               │
│ .               │
│ ifDescr.48 = B22  │
│ ifDescr.49 = B23  │
│ ifDescr.50 = B24  │
│ .               │
│ .               │
└ ─── ─── ─── ───┘

```

Figure 13-5. Matching Internal Port Numbers to External Slot/Port Numbers

Configuring Inbound Rate-Limiting for Broadcast and Multicast Traffic

Rate-limiting (throttling) of inbound broadcast and multicast traffic on the switch can be configured, which helps prevent the switch from being disrupted by traffic storms if they occur on the rate-limited port. The rate-limiting is implemented as a percentage of the total available bandwidth on the port.

The **rate-limit** command can be executed from the global or interface context, for example:

```

ProCurve(config)# interface 3 rate-limit bcst in
percent 10

or

ProCurve(config)# interface 3
ProCurve(eth-3)# rate-limit bcst in percent 10

```

Syntax: rate-limit < bcast | mcast > in percent <0-100>
no rate-limit <bcast | mcast> in

Enables rate-limiting and sets limits for the specified inbound broadcast or multicast traffic. Only the amount of traffic specified by the percent is forwarded.

Default: Disabled

For example, if you want to set a limit of 50 percent on inbound broadcast traffic for port 3, you can first enter interface context for port 3 and then execute the **rate-limit** command, as shown in Figure 13-6. Only 50 percent of the inbound broadcast traffic will be forwarded.

```
ProCurve(config)# int 3
ProCurve(eth-3)# rate-limit bcast in percent 50

ProCurve(eth-3)# show rate-limit bcast
Broadcast-Traffic Rate Limit Maximum %
```

Port	Inbound Limit	Mode	Radius Override
1	Disabled	Disabled	No-override
2	Disabled	Disabled	No-override
3	50	%	No-override
4	Disabled	Disabled	No-override
5	Disabled	Disabled	No-override

Figure 13-6. Example of Inbound Broadcast Rate-limiting of 50% on Port 3

If you rate-limit multicast traffic on the same port, the multicast limit is also in effect for that port, as shown in Figure 13-7. Only 20 percent of the multicast traffic will be forwarded.

```

ProCurve(eth-3)# rate-limit mcast in percent 20
ProCurve(eth-3)# show rate-limit mcast

Multicast-Traffic Rate Limit Maximum %

Port | Inbound Limit Mode      Radius Override
-----+-----
1    | Disabled      Disabled No-override
2    | Disabled      Disabled No-override
3    | 20            %      No-override
4    | Disabled      Disabled No-override

```

Figure 13-7. Example of Inbound Multicast Rate-limiting of 20% on Port 3

To disable rate-limiting for a port enter the **no** form of the command.

```

ProCurve(eth-3)# no rate-limit mcast in
ProCurve(eth-3)# show rate-limit mcast

Multicast-Traffic Rate Limit Maximum %

Port | Inbound Limit Mode      Radius Override
-----+-----
1    | Disabled      Disabled No-override
2    | Disabled      Disabled No-override
3    | Disabled      Disabled No-override
4    | Disabled      Disabled No-override

```

Figure 13-8. Example of Disabling Inbound Multicast Rate-limiting for Port 3

Operating Notes

- This rate-limiting option does not limit unicast traffic.
- This option does not include outbound multicast rate-limiting.

Guaranteed Minimum Bandwidth (GMB)

Feature	Default	Menu	CLI	Web
bandwidth-min output	Per-Queue: 2%-3%-30%-10% 10%-10%-15%-20%	n/a	page 13-25	n/a
show bandwidth output [<i>port-list</i>]	n/a	n/a	page 13-28	n/a

Introduction

Guaranteed Minimum Bandwidth (GMB) provides a method for ensuring that each of a given port's outbound traffic priority queues has a specified minimum consideration for sending traffic out on the link to another device. This can prevent a condition where applications generating lower-priority traffic in the network are frequently or continually "starved" by high volumes of higher-priority traffic. You can configure GMB per-port.

Terminology

Oversubscribed Queue: The condition where there is insufficient bandwidth allocated to a particular outbound priority queue for a given port. If additional, unused bandwidth is not available, the port delays or drops the excess traffic.

GMB Operation

The switch services per-port outbound traffic in a descending order of priority; that is, from the highest priority to the lowest priority. By default, each port offers eight prioritized, outbound traffic queues. Tagged VLAN traffic is prioritized according to the 802.1p priority the traffic carries. Untagged VLAN traffic is assigned a priority of "0" (normal).

Table 13-1. Per-Port Outbound Priority Queues

802.1p Priority Settings in Tagged VLAN Packets*	Outbound Priority Queue for a Given Port
1 (low)	1
2 (low)	2
0 (normal)	3
3 (normal)	4
4 (medium)	5
5 (medium)	6
6 (high)	7
7 (high)	8

*The switch processes outbound traffic from an untagged port at the "0" (normal) priority level.

You can use GMB to reserve a specific percentage of each port's available outbound bandwidth for each of the eight priority queues. This means that regardless of the amount of high priority outbound traffic on a port, you can ensure that there will always be bandwidth reserved for lower-priority traffic.

Since the switch services outbound traffic according to priority (highest to lowest), the highest-priority outbound traffic on a given port automatically receives the first priority in servicing. Thus, in most applications, it is necessary only to specify the minimum bandwidth you want to allocate to the lower priority queues. In this case, the high-priority traffic automatically receives all unassigned bandwidth without starving the lower-priority queues.

Conversely, configuring a bandwidth minimum on only the high-priority outbound queue of a port (and not providing a bandwidth minimum for the lower-priority queues) is not recommended because it may "starve" the lower-priority queues. (See the **Note** on page 13-24.)

Note

For a given port, when the demand on one or more outbound queues exceeds the minimum bandwidth configured for those queues, the switch apportions unallocated bandwidth to these queues on a priority basis. As a result, specifying a minimum bandwidth for a high-priority queue but not specifying a minimum for lower-priority queues can starve the lower-priority queues during periods of high demand on the high priority queue. For example, if a port configured to allocate a minimum bandwidth of 80% for outbound high-priority traffic experiences a demand above this minimum, then this burst starves lower-priority queues that *do not have a minimum configured*. Normally, this will not altogether halt lower priority traffic on the network, but will likely cause delays in the delivery of the lower-priority traffic.

The sum of the GMB settings for all outbound queues on a given port cannot exceed 100%.

Impacts of QoS Queue Configuration on GMB Operation

The section on “*Configuring Guaranteed Minimum Bandwidth for Outbound Traffic*” assumes the ports on the switch offer eight prioritized, outbound traffic queues. This may not always be the case, however, since the switch supports a QoS queue configuration feature that allows you to reduce the number of outbound queues from eight (the default) to four queues or two.

Changing the number of queues affects the GMB commands (**interface bandwidth-min** and **show bandwidth output**) such that they operate only on the number of queues currently configured. If the queues are reconfigured, the guaranteed minimum bandwidth per queue is automatically re-allocated according to the following percentages:

Table 13-2. Default GMB Percentage Allocations per QoS Queue Configuration

802.1p Priority	8 Queues (default)	4 Queues	2 Queues
1 (lowest)	2%	10%	90%
2	3%		
0 (normal)	30%	70%	
3	10%		
4	10%	10%	10%
5	10%	10%	
6	15%		
7 (highest)	20%		

Note

For more information on queue configuration and the associated default minimum bandwidth settings, refer to the chapter titled “*Quality of Service (QoS): Managing Bandwidth More Effectively*” in the *Advanced Traffic Management Guide* for your switch.

Configuring Guaranteed Minimum Bandwidth for Outbound Traffic

For any port or group of ports you can configure either the default minimum bandwidth settings for each outbound priority queue or a customized bandwidth allocation. For most applications, ProCurve recommends configuring GMB with the same values on all ports on the switch so that the outbound traffic profile is consistent for all outbound traffic. However, there may be instances where it may be advantageous to configure special profiles on connections to servers or to the network infrastructure (such as links to routers, other switches, or to the network core).

Syntax: [no] int < port-list > bandwidth-min output

Configures the default minimum bandwidth allocation for the outbound priority queue for each port in < port-list >. The default values per priority queue are:

- Queue 1 (low priority): 2%
- Queue 2 (low priority): 3%
- Queue 3 (normal priority): 30%
- Queue 4 (normal priority): 10%
- Queue 5 (medium priority): 10%
- Queue 6 (medium priority): 10%
- Queue 7 (high priority): 15%
- Queue 8 (high priority): 20%

*The **no** form of the command disables GMB for all ports in < port-list >. In this state, which is the equivalent of setting all outbound queues on a port to **0** (zero), a high level of higher-priority traffic can starve lower-priority queues, which can slow or halt lower-priority traffic in the network. You can configure bandwidth minimums from either the global configuration level (as shown above) or from the port context level. For information on outbound port queues, refer to table 13-1, “Per-Port Outbound Priority Queues” on page 13-23.*

Syntax: [no] int < port-list > bandwidth-min output

[< queue1% > < queue2% > < queue3% > < queue4% > < queue5% >
< queue6% > < queue7% > < queue8% >]

For ports in < port-list >, specifies the minimum outbound bandwidth as a percent of the total bandwidth for each outbound queue. The queues receive service in descending order of priority. You must specify a bandwidth percent value for all eight queues, and the sum of the bandwidth percentages must not exceed 100%. (0 is a value for a queue percentage setting.) Configuring a total of less than 100% across the eight queues results in unallocated bandwidth that remains harmlessly unused unless a given queue becomes oversubscribed. In this case, the unallocated bandwidth is apportioned to oversubscribed queues in descending order of priority. For example, if you configure a minimum of 10% for queues 1 - 7, and 0% for queue 8, then the unallocated bandwidth will be available to all eight queues in the following prioritized order:

1. Queue 8 (high priority)
2. Queue 7 (high priority)
3. Queue 6 (medium priority)
4. Queue 5 (medium priority)
5. Queue 4 (normal priority)
6. Queue 3 (normal priority)
7. Queue 2 (low priority)
8. Queue 1 (low priority)

A setting of 0 (zero %) on a queue means that no bandwidth minimum is specifically reserved for that queue for each of the ports in < port-list >. Also, there is no benefit to setting the high-priority queue (queue 8) to 0 (zero) unless you want the medium queue (queue 4) to be able to support traffic bursts above its guaranteed minimum.

(continued)

Notes: *Configuring 0% for a queue can result in that queue being starved if any higher queue becomes over-subscribed and is then given all unused bandwidth.*

The switch applies the bandwidth calculation to the link speed the port is currently using. For example, if a 10/100 Mbs port negotiates to 10 Mbps on the link, then it bases its GMB calculations on 10 Mbps; not 100 Mbps.

*Use **show bandwidth output < port-list >** to display the current GMB configuration. (The **show config** and **show running** commands do not include GMB configuration data.)*

For example, suppose you wanted to configure the following outbound minimum bandwidth availability for ports A1 and A2:

Priority of Outbound Port Queue	Minimum Bandwidth %	Effect on Outbound Bandwidth Allocation
8	20	Queue 8 has the first priority use of all outbound bandwidth not specifically allocated to queues 1 - 7. If, for example, bandwidth allocated to queue 5 is not being used and queues 7 and 8 become oversubscribed, queue 8 has first-priority use of the unused bandwidth allocated to queue 5.
7	15	Queue 7 has a guaranteed minimum bandwidth of 15% available for outbound traffic. If queue 7 becomes oversubscribed and queue 8 is not already using all of the unallocated bandwidth, then queue 7 can use the unallocated bandwidth. Also, any unused bandwidth allocated to queues 6 to queue 1 is available to queue 7 if queue 8 has not already claimed it.
6	10	Queue 6 has a guaranteed minimum bandwidth of 10% and, if oversubscribed, is subordinate to queues 8 and 7 in priority for any unused outbound bandwidth available on the port.
5	10	Queue 5 has a guaranteed minimum bandwidth of 10% and, if oversubscribed, is subordinate to queues 8, 7, and 6 for any unused outbound bandwidth available on the port.
4	10%	Queue 4 has a guaranteed minimum bandwidth of 10% and, if oversubscribed, is subordinate to queues, 8, 7, 6, and 5 for any unused outbound bandwidth available on the port.
3	30%	Queue 3 has a guaranteed minimum bandwidth of 30% and, if oversubscribed, is subordinate to queues, 8, 7, 6, 5, and 4 for any unused outbound bandwidth available on the port.
2	3%	Queue 2 has a guaranteed minimum bandwidth of 3% and, if oversubscribed, is subordinate to queues, 8, 7, 6, 5, 4, and 3 for any unused outbound bandwidth available on the port.
1	2%	Queue 1 has a guaranteed minimum bandwidth of 2% and, if oversubscribed, is subordinate to all the other queues for any unused outbound bandwidth available on the port.

Either of the following commands configures ports A1 through A5 with bandwidth settings:

```
ProCurve(config)#int a1-a5 bandwidth-min output 2 3 30 10  
10 10 15 20
```

```
ProCurve(eth-A1-A5)#bandwidth-min output 2 3 30 10 10 10  
15 20
```

Displaying the Current Guaranteed Minimum Bandwidth Configuration

This command displays the per-port GMB configuration in the running-config file.

Syntax: show bandwidth output [*port-list*]

*Without [port-list], this command lists the GMB configuration for all ports on the switch. With [port-list], this command lists the GMB configuration for the specified ports. This command operates the same way in any CLI context. If the command lists **Disabled** for a port, there are no bandwidth minimums configured for any queue on the port. (Refer to the description of the **no form of the bandwidth-min output command on page 13-25.**)*

For example, to display the GMB configuration resulting from either of the above commands:

```
ProCurve(config)# show bandwidth output a1-a5
```

Outbound Guaranteed Minimum Bandwidth %									
Port	Q1	Q2	Q3	Q4	Q5	Q6	Q7	Q8	
A1	2	3	30	10	10	10	15	20	
A2	2	3	30	10	10	10	15	20	
A3	2	3	30	10	10	10	15	20	
A4	2	3	30	10	10	10	15	20	
A5	2	3	30	10	10	10	15	20	

User-Configured Minimum Bandwidth Settings

Figure 13-9. Example of Listing the Guaranteed Minimum Bandwidth Configuration

This is how the preceding listing of the GMB configuration would appear in the startup-config file.

```
ProCurve(config)# show config status
Running configuration is same as the startup configuration.
ProCurve(config)# show config

Startup configuration:

; J8697A Configuration Editor; Created on release #K.11.00

hostname "ProCurve"
module 1 type J8697A
snmp-server community "public" Unrestricted
vlan 1
  name "DEFAULT_VLAN"
  untagged A1-A24
  ip address dhcp-bootp
  exit
interface A1
  bandwidth-min output 2 3 30 10 10 10 15 20
  exit
interface A2
  bandwidth-min output 2 3 30 10 10 10 15 20
  exit
interface A3
  bandwidth-min output 2 3 30 10 10 10 15 20
  exit
interface A4
  bandwidth-min output 2 3 30 10 10 10 15 20
  exit
interface A5
  bandwidth-min output 2 3 30 10 10 10 15 20
  exit
```

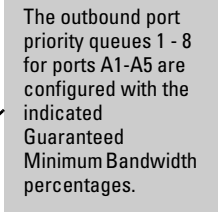


Figure 13-10. Example of GMB Settings Listed in the “show config” Output

GMB Operating Notes

Impact of QoS Queue Configuration on GMB commands. Changing the number of queues affects the GMB commands (**interface bandwidth-min** and **show bandwidth output**) to operate only on the number of queues currently configured. In addition, when the **qos queue-config** command is executed, any previously configured **bandwidth-min output** settings are removed from the startup configuration. Refer to Table 13-2 on page 13-24 for the default GMB percentage allocations per number of queues.

Jumbo Frames

Feature	Default	Menu	CLI	Web
display VLAN jumbo status	n/a	—	13-33	—
configure jumbo VLANs	Disabled	—	13-35	—

The *Maximum Transmission Unit* (MTU) is the maximum size IP frame the switch can receive for Layer 2 frames inbound on a port. The switch drops any inbound frames larger than the MTU allowed on the port. Ports operating at a minimum of 10 Mbps on the ProCurve 3500 switches and 1 Gbps on the other switches covered in this guide can accept forward frames of up to 9220 bytes (including four bytes for a VLAN tag) when configured for jumbo traffic. You can enable inbound jumbo frames on a per-VLAN basis. That is, on a VLAN configured for jumbo traffic, all ports belonging to that VLAN and *operating* at a minimum of 10 Mbps on the ProCurve 3500 switches and 1 Gbps on the other switches covered in this guide allow inbound jumbo frames of up to 9220 bytes.

Switch Model	Minimum Speed for Jumbo Traffic
3500	10 Mbps
All others in this guide	1 Gbps

Terminology

Jumbo Frame: An IP frame exceeding 1522 bytes in size. The maximum Jumbo frame size is 9220 bytes. (This size includes 4 bytes for the VLAN tag.)

Jumbo VLAN: A VLAN configured to allow inbound jumbo traffic. All ports belonging to a jumbo and operating at 1 Gbps or higher can receive jumbo frames from external devices. If the switch is in a meshed domain, then all meshed ports (operating at 1 Gbps or higher) on the switch will accept jumbo traffic from other devices in the mesh.

MTU (*Maximum Transmission Unit*): This is the maximum-size IP frame the switch can receive for Layer 2 frames inbound on a port. The switch allows jumbo frames of up to 9220 bytes.

Standard MTU: An IP frame of 1522 bytes in size. (This size includes 4 bytes for the VLAN tag.)

Operating Rules

- **Required Port Speed:** This feature allows inbound and outbound jumbo frames on ports operating at a minimum of 10 Mbps on the ProCurve 3500 switches and 1 Gbps on the other switches covered in this guide.
- **Switch Meshing:** If you enable jumbo traffic on a VLAN, then all meshed ports on the switch will be enabled to support jumbo traffic. (On a given meshed switch, every meshed port operating at 1 Gbps or higher becomes a member of every VLAN configured on the switch.)
- **GVRP Operation:** A VLAN enabled for jumbo traffic cannot be used to create a dynamic VLAN. A port belonging to a statically configured, jumbo-enabled VLAN cannot join a dynamic VLAN.
- **Port Adds and Moves:** If you add a port to a VLAN that is already configured for jumbo traffic, the switch enables that port to receive jumbo traffic. If you remove a port from a jumbo-enabled VLAN, the switch disables jumbo traffic capability on the port only if the port is not currently a member of another jumbo-enabled VLAN. This same operation applies to port trunks.
- **Jumbo Traffic Sources:** A port belonging to a jumbo-enabled VLAN can receive inbound jumbo frames through any VLAN to which it belongs, including non-jumbo VLANs. For example, if VLAN 10 (without jumbos enabled) and VLAN 20 (with jumbos enabled) are both configured on a switch, and port 1 belongs to both VLANs, then port 1 can receive jumbo traffic from devices on either VLAN. For a method to allow only some ports in a VLAN to receive jumbo traffic, refer to “Configuring a Maximum Frame Size” on page 13-35.

Configuring Jumbo Frame Operation

Command	Page
show vlans	13-33
show vlans ports < port-list >	13-34
show vlans < vid >	13-35
jumbo	13-35
jumbo max-frame-size	13-35

Overview

1. Determine the VLAN membership of the ports or trunks through which you want the switch to accept inbound jumbo traffic. For operation with GVRP enabled, refer to the GVRP topic under “Operating Rules”, above.
2. Ensure that the ports through which you want the switch to receive jumbo frames are operating at least at gigabit speed. (Check the **Mode** field in the output for the **show interfaces brief < port-list >** command.)
3. Use the **jumbo** command to enable jumbo frames on one or more VLANs statically configured in the switch. (All ports belonging to a jumbo-enabled VLAN can receive jumbo frames.)
4. Execute **write memory** to save your configuration changes to the startup-config file.

Viewing the Current Jumbo Configuration

Syntax: show vlans

*Lists the static VLANs configured on the switch and includes a **Jumbo** column to indicate which VLANs are configured to support inbound jumbo traffic. All ports belonging to a jumbo-enabled VLAN can receive jumbo traffic. (For more information refer to “Configuring a Maximum Frame Size” on page 13-35.) See Figure 13-11, below.*

```
ProCurve(config)# show vlans
```

Status and Counters - VLAN Information

Maximum VLANs to support : 8
Primary VLAN : DEFAULT_VLAN
Management VLAN :

802.1Q	VLAN ID	Name	Status	Voice	Jumbo
1		DEFAULT_VLAN	Port-based	No	Yes
5		VLAN5	Port-based	No	No
22		VLAN22	Port-based	No	No

Indicates which static VLANs are configured to enable jumbo frames.

Figure 13-11. Example Listing of Static VLANs To Show Jumbo Status Per VLAN

Syntax: show vlans ports < port-list >

*Lists the static VLANs to which the specified port(s) belong, including the **Jumbo** column to indicate which VLANs are configured to support jumbo traffic. Entering only one port in < port-list > results in a list of all VLANs to which that port belongs. Entering multiple ports in < port-list > results in a superset list that includes the VLAN memberships of all ports in the list, even though the individual ports in the list may belong to different subsets of the complete VLAN listing. For example, if port 1 belongs to VLAN 1, port 2 belongs to VLAN 10, and port 3 belongs to VLAN 15, then executing this command with a < port-list > of **1-3** results in a listing of all three VLANs, even though none of the ports belong to all three VLANs. (Refer to Figure 13-12.)*

```
ProCurve# show vlans ports 1-3
```

Status and Counters - VLAN Information - for ports 1-3

802.1Q VLAN ID	Name	Status	Voice	Jumbo
1	DEFAULT_VLAN	Port-based	No	Yes
10	VLAN10	Port-based	No	No
15	VLAN15	Port-based	No	No

Indicates which static VLANs are configured to enable jumbo frames.

Figure 13-12. Example of Listing the VLAN Memberships for a Range of Ports

Syntax: show vlans < vid >

This command shows port membership and jumbo configuration for the specified < vid >.

```
ProCurve(config)# show vlan 100
```

Status and Counters - VLAN Information - Ports - VLAN 100

802.1Q VLAN ID : 100
Name : VLAN100
Status : Port-based
Voice : No
Jumbo : No

Port	Information Mode	Unknown	VLAN	Status
1	Tagged	Learn		Up
2	Tagged	Learn		Up
3	Tagged	Learn		Up
4	Tagged	Learn		Down
5	Tagged	Learn		Up

Lists the ports belonging to VLAN 100 and whether the VLAN is enabled for jumbo frame traffic.

Figure 13-13. Example of Listing the Port Membership and Jumbo Status for a VLAN

Enabling or Disabling Jumbo Traffic on a VLAN

Syntax: vlan < vid > jumbo
[no] vlan < vid > jumbo

*Configures the specified VLAN to allow jumbo frames on all ports on the switch that belong to that VLAN. If the VLAN is not already configured on the switch, **vlan < vid > jumbo** also creates the VLAN. Note that a port belonging to one jumbo VLAN can receive jumbo frames through any other VLAN statically configured on the switch, regardless of whether the other VLAN is enabled for jumbo frames. The **[no]** form of the command disables inbound jumbo traffic on all ports in the specified VLAN that do not also belong to another VLAN that is enabled for jumbo traffic. In a VLAN context, the command forms are **jumbo** and **no jumbo**. (Default: Jumbos disabled on the specified VLAN.)*

Configuring a Maximum Frame Size

You can globally set a maximum frame size for Jumbo frames that will support values from 1518 bytes to 9216 bytes for untagged frames.

Syntax: jumbo max-frame-size <size>

Sets the maximum frame size for Jumbo frames. The range is from 1518 bytes to 9216 bytes.

Note: The **jumbo max-frame-size** is set on a **GLOBAL** level.

Default: 9216 bytes

Configuring IP MTU

Note

The following feature is available on the switches covered in this guide. Jumbos support is required. On switches that do not support this command, the IP MTU value is derived from the maximum frame size and is not configurable.

You can set the IP MTU globally by entering this command. The value of **max-frame-size** must be greater than or equal to 18 bytes more than the value selected for **ip-mtu**. For example, if **ip-mtu** is set to 8964, the **max-frame-size** is configured as 8982.

Syntax: jumbo ip-mtu <size>

Globally sets the IP MTU size. Values range between 1500 and 9198 bytes. This value must be 18 bytes less than the value of max-frame-size.

Default: 9198 bytes

SNMP Implementation

Jumbo Maximum Frame Size.

The maximum frame size for Jumbos is supported with the following proprietary MIB object:

hpSwitchMaxFrameSize OBJECT-TYPE

This is the value of the global **max-frame-size** supported by the switch. The default value is set to 9216 bytes.

Jumbo IP MTU.

The IP MTU for Jumbos is supported with the following proprietary MIB object:

hpSwitchIpMTU OBJECT-TYPE

This is the value of the global Jumbos IP MTU (or L3 MTU) supported by the switch. The default value is set to 9198 bytes (a value that is 18 bytes less than the largest possible maximum frame size of 9216 bytes). This object can only be used in switches which support **max-frame-size** and **ip-mtu** configuration.

Displaying the Maximum Frame Size

Use the **show jumbos** command to display the globally configured untagged maximum frame size for the switch.

```
ProCurve(config)# show jumbos

Jumbos Global Values

Configured : MaxFrameSize : 9216      Ip-MTU : 9198
In Use     : MaxFrameSize : 9216      Ip-MTU : 9198
```

Figure 14. Displaying the Maximum Frame Size and IP MTU Values

Operating Notes for Maximum Frame Size

- When you set a maximum frame size for Jumbo frames, it must be on a global level. You cannot use the **jumbo max-frame-size** command on a per-port or per-VLAN basis.
- The original way to configure Jumbo frames remains the same, which is per-VLAN, but you cannot set a maximum frame size per-VLAN.
- Jumbo support must be enabled for a VLAN from the CLI or through SNMP.
- Setting the maximum frame size does not require a reboot.
- When you upgrade to a version of software that supports setting the maximum frame size from a version that did not, the **max-frame-size** value is set automatically to 9216 bytes.
- Configuring a Jumbo maximum frame size on a VLAN allows frames up to **max-frame-size** even though other VLANs of which the port is a member are not enabled for Jumbo support.

Operating Notes for Jumbo Traffic-Handling

- ProCurve does not recommend configuring a voice VLAN to accept jumbo frames. Voice VLAN frames are typically small, and allowing a voice VLAN to accept jumbo frame traffic can degrade the voice transmission performance.
- You can configure the default, primary, and/or (if configured) the management VLAN to accept jumbo frames on all ports belonging to the VLAN.

- When the switch applies the default MTU (1522-bytes) to a VLAN, all ports in the VLAN can receive incoming frames of up to 1522 bytes in length. When the switch applies the jumbo MTU (9220 bytes) to a VLAN, all ports in that VLAN can receive incoming frames of up to 9220 bytes in length. A port receiving frames exceeding the applicable MTU drops such frames, causing the switch to generate an Event Log message and increment the “Giant Rx” counter (displayed by **show interfaces < port-list >**).
- The switch allows flow control and jumbo frame capability to co-exist on a port.
- The default MTU is 1522 bytes (including 4 bytes for the VLAN tag). The jumbo MTU is 9220 bytes (including 4 bytes for the VLAN tag).
- When a port is not a member of any jumbo-enabled VLAN, it drops all jumbo traffic. If the port is receiving “excessive” inbound jumbo traffic, the port generates an Event Log message to notify you of this condition. This same condition generates a Fault-Finder message in the Alert log of the switch’s web browser interface, and also increments the switch’s “Giant Rx” counter.
- If you do not want all ports in a given VLAN to accept jumbo frames, you can consider creating one or more jumbo VLANs with a membership comprised of only the ports you want to receive jumbo traffic. Because a port belonging to one jumbo-enabled VLAN can receive jumbo frames through any VLAN to which it belongs, this method enables you to include both jumbo-enabled and non-jumbo ports within the same VLAN. For example, suppose you wanted to allow inbound jumbo frames only on ports 6, 7, 12, and 13. However, these ports are spread across VLAN 100 and VLAN 200, and also share these VLANs with other ports you want excluded from jumbo traffic. A solution is to create a third VLAN with the sole purpose of enabling jumbo traffic on the desired ports, while leaving the other ports on the switch disabled for jumbo traffic. That is:

	VLAN 100	VLAN 200	VLAN 300
Ports	6-10	11-15	6, 7, 12, and 13
Jumbo-Enabled?	No	No	Yes

If there are security concerns with grouping the ports as shown for VLAN 300, you can either use source-port filtering to block unwanted traffic paths or create separate jumbo VLANs, one for ports 6 and 7, and another for ports 12 and 13.

- **Outbound Jumbo Traffic.** Any port operating at 1 Gbps or higher can transmit outbound jumbo frames through any VLAN, regardless of the jumbo configuration. The VLAN is not required to be jumbo-enabled, and the port is not required to belong to any other, jumbo enabled VLANs. This

can occur in situations where a non-jumbo VLAN includes some ports that do not belong to another, jumbo-enabled VLAN and some ports that do belong to another, jumbo-enabled VLAN. In this case, ports capable of receiving jumbo frames can forward them to the ports in the VLAN that do not have jumbo capability.




Figure 13-14. Forwarding Jumbo Frames Through Non-Jumbo Ports

Jumbo frames can also be forwarded out non-jumbo ports when the jumbo frames received inbound on a jumbo-enabled VLAN are routed to another, non-jumbo VLAN for outbound transmission on ports that have no memberships in other, jumbo-capable VLANs. Where either of the above scenarios is a possibility, the downstream device must be configured to accept the jumbo traffic. Otherwise, this traffic will be dropped by the downstream device.

- **Jumbo Traffic in a Switch Mesh Domain.** Note that if a switch belongs to a meshed domain, but does not have any VLANs configured to support jumbo traffic, then the meshed ports on that switch will drop any jumbo frames they receive from other devices. In this regard, if a mesh domain includes any ProCurve 1600M/2400M/2424M/4000M/8000M switches along with the switches covered in this guide configured to support jumbo traffic, only the switches covered in this guide will receive jumbo frames. The other switch models in the mesh will drop such frames. For more information on switch meshing, refer to the chapter titled “Switch Meshing” in the *Advanced Traffic Management Guide* for your switch.

Troubleshooting

A VLAN is configured to allow jumbo frames, but one or more ports drops all inbound jumbo frames. The port may not be operating at a minimum of 10 Mbps on the ProCurve 3500 switches or 1 Gbps on the other switches covered in this guide. Regardless of a port's configuration, if it is actually operating at a speed lower than 10 Mbps for ProCurve 3500 switches or 1 Gbps for the other switches, it drops inbound jumbo frames. For example, if a port is configured for **Auto** mode (**speed-duplex auto**), but has negotiated a 7 Mbps speed with the device at the other end of the link, then the port cannot receive inbound jumbo frames. To determine the actual operating speed of one or more ports, view the **Mode** field in the output for the following command:

```
show interfaces brief <port-list>
```

A non-jumbo port is generating “Excessive undersize/giant frames” messages in the Event Log. The switches can transmit outbound jumbo traffic on any port, regardless of whether the port belongs to a jumbo VLAN. In this case, another port in the same VLAN on the switch may be jumbo-enabled through membership in a different, jumbo-enabled VLAN, and may be forwarding jumbo frames received on the jumbo VLAN to non-jumbo ports. Refer to “Outbound Jumbo Traffic” on page 13-38.

Configuring for Network Management Applications

Contents

Using SNMP Tools To Manage the Switch	14-2
Overview	14-2
SNMP Management Features	14-3
Configuring for SNMP version 1 and 2c Access to the Switch	14-3
Configuring for SNMP Version 3 Access to the Switch	14-4
SNMP Version 3 Commands	14-5
Enabling SNMPv3	14-6
SNMPv3 Users	14-6
Group Access Levels	14-10
SNMPv3 Communities	14-10
Menu: Viewing and Configuring non-SNMP version 3 Communities	14-12
CLI: Viewing and Configuring SNMP Community Names	14-14
SNMP Notifications	14-16
Supported Notifications	14-16
General Steps for Configuring SNMP Notifications	14-17
SNMPv1 and SNMPv2c Traps	14-18
Configuring an SNMP Trap Receiver	14-18
Enabling SNMPv2c Informs	14-20
Configuring SNMPv3 Notifications	14-22
Managing Network Security Notifications	14-25
Enabling Link-Change Traps	14-27
Configuring the Source IP Address for SNMP Notifications ..	14-28
Displaying SNMP Notification Configuration	14-30
Configuring Listening Mode	14-32
Advanced Management: RMON	14-33

CLI-Configured sFlow with Multiple Instances	14-33
Terminology	14-33
Configuring sFlow	14-34
Viewing sFlow Configuration and Status	14-34
LLDP (Link-Layer Discovery Protocol)	14-37
Terminology	14-38
General LLDP Operation	14-40
LLDP-MED	14-40
Packet Boundaries in a Network Topology	14-40
Configuration Options	14-41
Options for Reading LLDP Information Collected by the Switch ..	14-43
LLDP and LLDP-MED Standards Compatibility	14-43
LLDP Operating Rules	14-44
Configuring LLDP Operation	14-45
Viewing the Current Configuration	14-45
Configuring Global LLDP Packet Controls	14-47
Configuring SNMP Notification Support	14-51
Configuring Per-Port Transmit and Receive Modes	14-52
Configuring Basic LLDP Per-Port Advertisement Content	14-53
Configuring Support for Port Speed and Duplex Advertisements	14-55
LLDP-MED (Media-Endpoint-Discovery)	14-56
LLDP-MED Topology Change Notification	14-59
LLDP-MED Fast Start Control	14-61
Advertising Device Capability, Network Policy, PoE Status and Location Data	14-61
Configuring Location Data for LLDP-MED Devices	14-65
Displaying Advertisement Data	14-70
Displaying Switch Information Available for Outbound Advertisements	14-71
Displaying LLDP Statistics	14-75
LLDP Operating Notes	14-77
LLDP and CDP Data Management	14-79
LLDP and CDP Neighbor Data	14-79
CDP Operation and Commands	14-81

Using SNMP Tools To Manage the Switch

Overview

You can manage the switch via SNMP from a network management station running an application such as ProCurve Manager (PCM) or ProCurve Manager Plus (PCM+). For more on PCM and PCM+, visit the ProCurve Networking web site at:

www.procurve.com

Click on **products index** in the sidebar, then click on the appropriate link appearing under the **Network Management** heading.

This section includes:

- An overview of SNMP management for the switch
- Configuring the switches for:
 - SNMP Communities (page 14-11)
 - Trap Receivers and Authentication Traps (page 14-17)
- Information on advanced management through RMON Support (page 14-34)

To implement SNMP management, the switch must have an IP address, configured either manually or dynamically (using DHCP or Bootp). If multiple VLANs are configured, each VLAN interface should have its own IP address. For DHCP use with multiple VLANs, refer to the section titled “The Primary VLAN” in the “Static Virtual LANs (VLANs)” chapter of the *Advanced Traffic Management Guide* for your switch.

Note

If you use the switch’s Authorized IP Managers and Management VLAN features, ensure that the SNMP management station and/or the choice of switch port used for SNMP access to the switch are compatible with the access controls enforced by these features. Otherwise, SNMP access to the switch will be blocked. For more on Authorized IP Managers, refer to the *Access Security Guide* for your switch. (The latest version of this guide is available on the ProCurve Networking web site.) For information on the Management VLAN feature, refer to the section titled “The Secure Management VLAN” in the “Static Virtual LANs (VLANs)” chapter of the *Advanced Traffic Management Guide* for your switch.

SNMP Management Features

SNMP management features on the switch include:

- SNMP version 1, version 2c, or version 3 over IP
- Security via configuration of SNMP communities (page 14-11)
- Security via authentication and privacy for SNMP Version 3 access
- Event reporting via SNMP
 - Version 1 traps
 - RMON: groups 1, 2, 3, and 9
- ProCurve Manager/Plus support
- Flow sampling using sFlow
- Standard MIBs, such as the Bridge MIB (RFC 1493), Ethernet MAU MIB (RFC 1515), and others.

The switch SNMP agent also uses certain variables that are included in a Hewlett-Packard proprietary MIB (Management Information Base) file. If you are using HP OpenView, you can ensure that it is using the latest version of the MIB file by downloading the file to the OpenView database. To do so, go to the ProCurve Networking web site at:

www.procurve.com

Click on **software updates**, then **MIBs**.

Configuring for SNMP version 1 and 2c Access to the Switch

SNMP access requires an IP address and subnet mask configured on the switch. (Refer to “IP Configuration” on page 8-2.) If you are using DHCP/Bootp to configure the switch, ensure that the DHCP/Bootp process provides the IP address. (Refer to “DHCP/Bootp Operation” on page 8-12.)

Once an IP address has been configured, the main steps for configuring SNMP version 1 and version 2c access management features are:

1. Configure the appropriate SNMP communities. (Refer to “SNMPv3 Communities” on page 14-11.)
2. Configure the appropriate trap receivers. (Refer to “SNMP Notifications” on page 14-17.)

In some networks, authorized IP manager addresses are not used. In this case, all management stations using the correct community name may access the switch with the View and Access levels that have been set for that community.

If you want to restrict access to one or more specific nodes, you can use the switch's IP Authorized Manager feature. (Refer to the *Access Security Guide* for your switch.)

Caution

For ProCurve Manager (PCM) version 1.5 or earlier (or any TopTools version), deleting the “public” community disables some network management functions (such as traffic monitoring, SNMP trap generation, and threshold setting). If network management security is a concern, and you are using the above software versions, ProCurve recommends that you change the write access for the “public” community to “Restricted”.

Configuring for SNMP Version 3 Access to the Switch

SNMP version 3 (SNMPv3) access requires an IP address and subnet mask configured on the switch. (Refer to “IP Configuration” on page 8-2.) If you are using DHCP/Bootp to configure the switch, ensure that the DHCP/Bootp process provides the IP address. (See “DHCP/Bootp Operation” on page 8-12.)

Once an IP address has been configured, the main steps for configuring SNMP version 3 access management features are:

1. Enable SNMPv3 for operation on the switch (Refer to “SNMP Version 3 Commands” on page 14-6)
2. Configure the appropriate SNMP users (Refer to “SNMPv3 Users” on page 14-7)
3. Configure the appropriate SNMP communities. (Refer to “SNMPv3 Communities” on page 14-11.)
4. Configure the appropriate trap receivers. (Refer to “SNMP Notifications” on page 14-17.)

In some networks, authorized IP manager addresses are not used. In this case, all management stations using the correct User and community name may access the switch with the View and Access levels that have been set for that community. If you want to restrict access to one or more specific nodes, you can use the switch's IP Authorized Manager feature. (Refer to the *Access Security Guide* for your switch.)

SNMP Version 3 Commands

SNMP version 3 (SNMPv3) adds some new commands to the CLI for configuring SNMPv3 functions. To enable SMNPv3 operation on the switch, use the **snmpv3 enable** command. An initial user entry will be generated with MD5 authentication and DES privacy.

You may (optionally) restrict access to only SNMPv3 agents by using the **snmpv3 only** command. To restrict write-access to only SNMPv3 agents, use the **snmpv3 restricted-access** command.

Caution

Restricting access to only version 3 messages will make the community named “public” inaccessible to network management applications (such as auto-discovery, traffic monitoring, SNMP trap generation, and threshold setting) from operating in the switch.

Syntax: [no] snmpv3 enable

Enable and disable the switch for access from SNMPv3 agents. This includes the creation of the initial user record.

[no] snmpv3 only

Enables or disables restrictions to access from only SNMPv3 agents. When enabled, the switch will reject all non-SNMPv3 messages.

[no] snmpv3 restricted-access

Enables or disables restrictions from all non-SNMPv3 agents to read only access.

show snmpv3 enable

Displays the operating status of SNMPv3.

show snmpv3 only

Displays status of message reception of non-SNMPv3 messages.

show snmpv3 restricted-access

Displays status of write messages of non-SNMPv3 messages.

Enabling SNMPv3

The **snmpv3 enable** command allows the switch to:

- Receive SNMPv3 messages.
- Configure initial users.
- Restrict non-version 3 messages to “read only” (optional).

Figure 14-1 shows an example of how to use the **snmpv3 enable** command.

Note:
SNMP
Version 3
Initial Users

To create new users, most SNMPv3 management software requires an initial user record to clone. The initial user record can be downgraded and provided with fewer features, but not upgraded by adding new features. For this reason it is recommended that when you enable SNMPv3, you also create a second user with SHA authentication and DES privacy.

```
ProCurve (config)# snmpv3 enable
SNMPv3 Initialization process.
Creating user 'initial'
Authentication Protocol: MD5
Enter authentication password: *****
Privacy protocol is DES
Enter privacy password: *****

User 'initial' is created
Would you like to create a user that uses SHA? y
Enter user name: templateSHA
Authentication Protocol: SHA
Enter authentication password: *****
Privacy protocol is DES
Enter privacy password: *****

User creation is done.  SNMPv3 is now functional.
Would you like to restrict SNMPv1 and SNMPv2c messages to have read only
access (you can set this later by the command 'snmp restrict-access'): n
```

The diagram shows three callout boxes with arrows pointing to the terminal output. The first box, 'Enable SNMPv3', points to the 'snmpv3 enable' command. The second box, 'Create initial user models for SNMPv3 Management Applications', points to the 'Creating user 'initial'' section. The third box, 'Set restriction on non-SNMPv3 messages', points to the final question about restricting SNMPv1 and SNMPv2c messages.

Figure 14-1. Example of SNMP version 3 Enable Command

SNMPv3 Users

To use SNMPv3 on the switch, you must configure the users that will be assigned to different groups. To configure SNMP users on the switch:

1. Configure users in the User Table with the **snmpv3 user** command. To view the list of configured users, enter the **show snmpv3 user** command (see “Adding Users” on page 14-8).
2. Assign users to Security Groups based on their security model with the **snmpv3 group** command (see “Assigning Users to Groups” on page 14-10).

Caution

If you add an SNMPv3 user without authentication and/or privacy to a group that requires either feature, the user will not be able to access the switch. Ensure that you add a user with the appropriate security level to an existing security group.

Adding Users. To configure an SNMPv3 user, you must first add the user name to the list of known users with the **snmpv3 user** command.

```
ProCurve(config)# snmpv3 user NetworkAdmin
ProCurve(config)# snmpv3 user NetworkMgr auth md5 authpass priv privpass
ProCurve(config)# show snmpv3 user
```

Status and Counters - SNMP v3 Global Configuration Information

User Name	Auth. Protocol	Privacy Protocol
initial	MD5	CFB AES-128
NetworkAdmin	MD5	CBC-DES

Figure 14-2. Adding SNMPv3 Users and Displaying SNMPv3 Configuration

SNMPv3 User Commands

Syntax: [no] snmpv3 user <user_name>

Adds or deletes a user entry for SNMPv3. Authorization and privacy are optional, but to use privacy, you must use authorization. When you delete a user, only the <user_name> is required.

[auth <md5 | sha> <auth_pass>]

*With authorization, you can set either MD5 or SHA authentication. The authentication password <auth_pass> must be 6-32 characters in length and is mandatory when you configure authentication.
Default: None*

[priv <des | aes> <priv_pass>]

*With privacy, the switch supports DES (56-bit) and AES (128-bit) encryption. The privacy password <priv_pass> must be 6-32 characters in length and is mandatory when you configure privacy.
Default: DES*

Note: *Only AES 128-bit and DES 56-bit encryption are supported as privacy protocols. Other non-standard encryption algorithms, such as AES-172, AES-256, and 3-DES are not supported.*

Listing Users. To display the management stations configured to access the switch with SNMPv3 and view the authentication and privacy protocols that each station uses, enter the **show snmpv3 user** command.

Syntax: show snmpv3 user

This example displays information about the management stations configured on VLAN 1 to access the switch.

```
ProCurve# configure terminal
ProCurve(config)# vlan 1
ProCurve(vlan-1)# show snmpv3 user

Status and Counters - SNMPv3 Global Configuration Information

User Name          Auth. Protocol    Privacy Protocol
-----
initial            MD5               CFB AES-128
NetworkAdmin       MD5               CBC-DES
```

Assigning Users to Groups. Then you must set the group access level for the user by assigning the user to a group. This is done with the **snmpv3 group** command. For more details on the MIBs access for a given group refer to “Group Access Levels” on page 14-11.

The screenshot shows a configuration session on a ProCurve switch. Two commands are used to assign users to groups: `snmpv3 group operatornoauth user NetworkAdmin sec-model ver3` and `snmpv3 group managerpriv user NetworkMgr sec-model ver3`. A subsequent `show snmpv3 group` command displays the configuration. Annotations with arrows point to these commands and a specific row in the table below.

```

ProCurve (config)# snmpv3 group operatornoauth user NetworkAdmin sec-model ver3
ProCurve (config)# snmpv3 group managerpriv user NetworkMgr sec-model ver3
ProCurve (config)# show snmpv3 group
  
```

Status and Counters - SNMP v3 Global Configuration Information

Security Name	Security Model	Group Name
CommunityManagerReadOnly	ver1	ComManagerR
CommunityManagerReadWrite	ver1	ComManagerRW
CommunityOperatorReadOnly	ver1	ComOperatorRW
CommunityOperatorReadWrite	ver1	ComOperatorRW
CommunityManagerReadOnly	ver2c	ComManagerR
CommunityManagerReadWrite	ver2c	ComManagerRW
CommunityOperatorReadOnly	ver2c	ComOperatorRW
CommunityOperatorReadWrite	ver2c	ComOperatorRW
NetworkMgr	ver3	ManagerPriv
NetworkAdmin	ver3	OperatorNoAuth

Annotations in the image include:

- "Add NetworkAdmin to operator noauth group" pointing to the first configuration command.
- "Add NetworkMgr to managerpriv group" pointing to the second configuration command.
- "Pre-assigned groups for access by Version 2c and version 1 management applications" pointing to the ver1 and ver2c rows in the table.

Figure 14-3. Example of Assigning Users to Groups

SNMPv3 Group Commands

Syntax: [no] snmpv3 group

This command assigns or removes a user to a security group for access rights to the switch. To delete an entry, all of the following three parameters must be included in the command.

group <group_name>

This parameter identifies the group that has the privileges that will be assigned to the user. For more details refer to “Group Access Levels” on page 14-11.

user <user_name>

*This parameter identifies the user to be added to the access group. This must match the user name added with the **snmpv3 user** command.*

sec-model <ver1 | ver2c | ver3>

This defines which security model to use for the added user. A SNMPv3 access Group should only use the ver3 security model.

Group Access Levels

The switch supports eight predefined group access levels. There are four levels for use with version 3 users and four are used for access by version 2c or version 1 management applications.

Group Name	Group Access Type	Group Read View	Group Write View
managerpriv	Ver3 Must have Authentication and Privacy	ManagerReadView	ManagerWriteView
managerauth	Ver3 Must have Authentication	ManagerReadView	ManagerWriteView
operatorauth	Ver3 Must have Authentication	OperatorReadView	DiscoveryView
operatornoauth	Ver3 No Authentication	OperatorReadView	DiscoveryView
commanagerrw	Ver2c or Ver1	ManagerReadView	ManagerWriteView
commanagerr	Ver2c or Ver1	ManagerReadView	DiscoveryView
comoperatorrw	Ver2c or Ver1	OperatorReadView	OperatorReadView
comoperatorr	Ver2c or Ver1	OperatorReadView	DiscoveryView

Each view allows you to view or modify a different set of MIBs.

- **Manager Read View** – access to all managed objects
- **Manager Write View** – access to all managed objects *except* the following: vacmContextTable, vacmAccessTable, vacmViewTreeFamilyTable
- **OperatorReadView** – no access to icfSecurityMIB, hpSwitchIpTftp-Mode, vacmContextTable, vacmAccessTable, vacmViewTreeFamilyTable, usmUserTable, snmpCommunityTable
- **Discovery View** – Access limited to samplingProbe MIB.

Note

All access groups and views are predefined on the switch. There is no method to modify or add groups or views to those that are pre-defined on the switch.

SNMPv3 Communities

SNMP communities are supported by the switch to allow management applications that use version 2c or version 1 to access the switch. The communities are mapped to Group Access Levels that are used for version 2c or version 1 support. For more information refer to “Group Access Levels” on page 14-11. This mapping will happen automatically based on the communities access privileges, but special mappings can be added with the **snmpv3 community** command.

Syntax: [no] snmpv3 community

*This command maps or removes a mapping of a community name to a group access level. To remove a mapping you, only need to specify the **index_name** parameter.*

index <index_name>

This is an index number or title for the mapping. The values of 1-5 are reserved and can not be mapped.

name <community_name>

This is the community name that is being mapped to a group access level.

sec-name <security_name>

This is the group level to which the community is being mapped. For more information refer to “Group Access Levels” on page 14-11.

tag <tag_value>

This is used to specify which target address may have access by way of this index reference.

Figure 14-4 shows the assigning of the Operator community on MgrStation1 to the CommunityOperatorReadWrite group. Any other Operator only has an access level of CommunityOperatorReadOnly

```
ProCurve (config)# snmpv3 community index 30 name Operator sec-name
CommunityManagerReadWrite tag MgrStation1
ProCurve (config)# show snmpv3 community
```

Index Name	Community Name	Security Name
1	public	CommunityManagerReadWrite
2	Operator	CommunityOperatorReadOnly
3	Manager	CommunityManagerReadWrite
30	Operator	CommunityManagerReadWrite

Figure 14-4. Assigning a Community to a Group Access Level

SNMP Community Features

Feature	Default	Menu	CLI	Web
show SNMP communities	n/a	page 14-13	page 14-15	—
configure identity information	none	—	page 14-16	—
configure community names	public	page 14-13	page 14-16	—
MIB view for a community name (operator, manager)	manager	"	"	—
write access for default community name	unrestricted	"	"	—

Use SNMP communities to restrict access to the switch by SNMP management stations by adding, editing, or deleting SNMP communities. You can configure up to five SNMP communities, each with either an operator-level or a manager-level view, and either restricted or unrestricted write access.

Using SNMP requires that the switch have an IP address and subnet mask compatible with your network.

Caution

For ProCurve Manager (PCM) version 1.5 or earlier (or any TopTools version), deleting the “public” community disables some network management functions (such as traffic monitoring, SNMP trap generation, and threshold setting). If network management security is a concern, and you are using the above software versions, ProCurve recommends that you change the write access for the “public” community to “Restricted”.

Menu: Viewing and Configuring non-SNMP version 3 Communities

To View, Edit, or Add SNMP Communities:

1. From the Main Menu, Select:
 2. **Switch Configuration...**
 6. **SNMP Community Names**

Configuring for Network Management Applications
Using SNMP Tools To Manage the Switch

Note: This screen gives an overview of the SNMP communities that are currently configured. All fields in this screen are read-only.

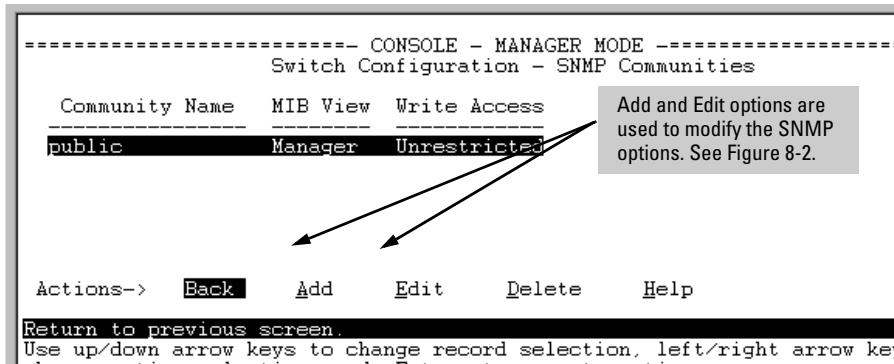


Figure 14-5. The SNMP Communities Screen (Default Values)

2. Press **[A]** (for **Add**) to display the following screen:

If you are adding a community, the fields in this screen are blank.
 If you are editing an existing community, the values for the currently selected Community appear in the fields.

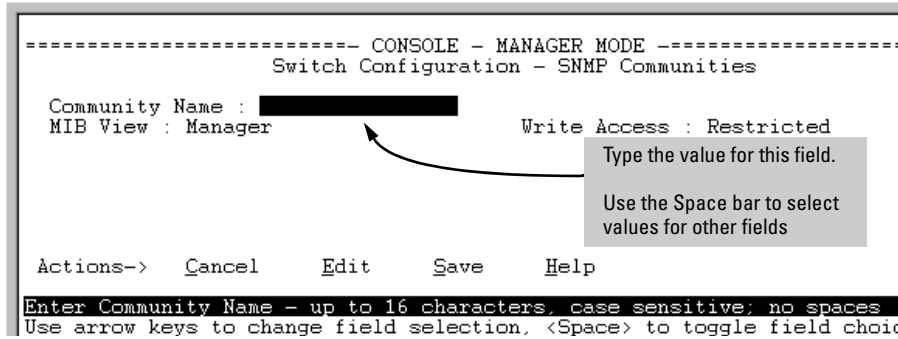


Figure 14-6. The SNMP Add or Edit Screen

Need Help? If you need information on the options in each field, press **[Enter]** to move the cursor to the Actions line, then select the **H**elp option on the Actions line. When you are finished with Help, press **[E]** (for **E**dit) to return the cursor to the parameter fields.

3. Enter the name you want in the Community Name field, and use the Space bar to select the appropriate value in each of the other fields. (Use the **[Tab]** key to move from one field to the next.)
4. Press **[Enter]**, then **[S]** (for **S**ave).

CLI: Viewing and Configuring SNMP Community Names

Community Name Commands	Page
show snmp-server [<i><community-string></i>]	14-15
[no] snmp-server	14-16
[community <i><community-str></i>]	14-16
[host <i><community-str></i> <i><ip-addr></i>] [<i><none debug all not-info critical></i>]	14-19
[enable traps <i><authentication></i>]	14-27
[enable traps link-change <i><port-list></i>]	14-28

Listing Community Names and Values. This command lists the data for currently configured SNMP community names (along with trap receivers and the setting for authentication traps — refer to “SNMP Notifications” on page 14-17).

Syntax: show snmp-server [*<community-string>*]

This example lists the data for all communities in a switch; that is, both the default “public” community name and another community named “blue-team”

```

ProCurve# show snmp-server

SNMP Communities

Community Name  MIB View  Write Access
-----
public          Manager   Unrestricted
blue-team       Operator  Restricted

Trap Receivers

Send Authentication Traps [No] : No

Address          Community  Events Sent in Trap
-----

```

Figure 14-7. Example of the SNMP Community Listing with Two Communities

To list the data for only one community, such as the “public” community, use the above command with the community name included. For example:

```
ProCurve# show snmp-server public
```

Configuring Community Names and Values. The **snmp-server** command enables you to add SNMP communities with either default or specific access attributes, and to delete specific communities.

Syntax: [no] snmp-server community < community-name >

Configures a new community name. If you do not also specify **operator** or **manager**, the switch automatically assigns the community to the **operator** MIB view. If you do not specify **restricted** or **unrestricted**, the switch automatically assigns the community to **restricted** (read-only) access. The **no** form uses only the < **community-name** > variable and deletes the named community from the switch.

[operator | manager]

*Optionally assigns an access level. At the **operator** level the community can access all MIB objects except the CONFIG MIB. At the **manager** level the community can access all MIB objects.*

[restricted | unrestricted]

*Optionally assigns MIB access type. Assigning the **restricted** type allows the community to read MIB variables, but not to set them. Assigning the **unrestricted** type allows the community to read and set MIB variables.*

For example, to add the following communities:

Community	Access Level	Type of Access
red-team	manager <i>(Access to all MIB objects.)</i>	unrestricted <i>(read/write)</i>
blue-team	operator <i>(Access to all MIB objects except the CONFIG MIB.)</i>	restricted <i>(read-only)</i>

```
ProCurve(config)# snmp-server community red-team  
manager unrestricted  
ProCurve(config)# snmp-server community blue-team  
operator restricted
```

To eliminate a previously configured community named "gold-team":

```
ProCurve(config) # no snmp-server community gold-team
```


SNMP Notifications

The switches covered in this guide support:

- SNMP version 1 or SNMP version 2c traps
- SNMPv2c informs
- SNMPv3 notification process, including traps

This section describes how to configure a switch to send network security and link-change notifications to configured trap receivers.

Supported Notifications

By default, the following notifications are enabled on a switch:

- Manager password changes
- SNMP authentication failure
- Link-change traps: when the link on a port changes from up to down (linkDown) or down to up (linkUp)
- Port-security (web, MAC, or 802.1X) authentication failure
- Invalid password entered in a login attempt through a direct serial, Telnet, or SSH connection
- Inability to establish a connection with the RADIUS or TACACS+ authentication server
- DHCP snooping events
- ARP protection events

In addition, you can enable the switch to send the following types of notifications to configured trap receivers. For information on how to configure each notification, refer to the ProCurve software guide under which the notification is listed.

- *Management and Configuration Guide:*
 - Configuration changes
 - ICMP rate-limiting
 - Instrumentation monitoring
 - Link-Layer Discovery Protocol (LLDP)
 - Ping tests
 - Power over Ethernet (POE): port toggle, power limit
 - RMON

- *Advance Traffic Management Guide:*
 - Loop protection
 - Spanning Tree (STP, RSTP, MSTP)
- *Access Security Guide:*
 - MAC lockdown
 - MAC lockout
 - Uni-Directional Link Detection (UDLD)
 - Virus throttling
- *Multicast and Routing Guide:*
 - OSPF
 - PIM
 - Virtual Router Redundancy Protocol (VRRP)

General Steps for Configuring SNMP Notifications

To configure SNMP notifications, follow these general steps:

1. Determine the versions of SNMP notifications that you want to use in your network.

If you want to use SNMPv1 and SNMPv2c traps, you must also configure a trap receiver. Refer to the following sections and follow the required configuration procedures:

- “SNMPv1 and SNMPv2c Traps” on page 14-19
- “Configuring an SNMP Trap Receiver” on page 14-19
- “Enabling SNMPv2c Informs” on page 14-21

If you want to use SNMPv3 notifications (including traps), you must also configure an SNMPv3 management station. Follow the required configuration procedure in the following section:

- “Configuring SNMPv3 Notifications” on page 14-23

2. To reconfigure any of the SNMP notifications that are enabled by default to be sent to a management station (trap receiver), refer to these sections:

-

- “Enabling Link-Change Traps” on page 14-28

3. (Optional) Refer to the following sections to configure optional SNMP notification features and verify the current configuration:

- “Configuring the Source IP Address for SNMP Notifications” on page 14-29
- “Displaying SNMP Notification Configuration” on page 14-31

SNMPv1 and SNMPv2c Traps

The switches covered in this guide support the following functionality from earlier SNMP versions (SNMPv1 and SNMPv2c):

- **Trap receivers:** A *trap receiver* is a management station to which the switch sends SNMP traps and (optionally) event log messages sent from the switch. From the CLI you can configure up to ten SNMP trap receivers to receive SNMP traps from the switch.
- **Fixed or “Well-Known” Traps:** A switch automatically sends fixed traps (such as “coldStart”, “warmStart”, “linkDown”, and “linkUp”) to trap receivers using the **public** community name. These traps cannot be redirected to other communities. If you change or delete the default **public** community name, these traps are not sent.
- **Thresholds:** A switch automatically sends all messages created when a system threshold is reached to the network management station that configured the threshold, regardless of the trap receiver configuration.

Configuring an SNMP Trap Receiver

Use the **snmp-server host** command to configure a trap receiver that can receive SNMPv1 and SNMPv2c traps, and (optionally) event log messages. When you configure a trap receiver, you specify its community membership, management station IP address, and (optionally) the type of event log messages to be sent.

If you specify a community name that does not exist—that is, has not yet been configured on the switch—the switch still accepts the trap receiver assignment. However, no traps will be sent to that trap receiver until the community to which it belongs has been configured on the switch.

Syntax: snmp-server host <ipv4-addr | ipv6-addr> <community name>

*Configures a destination network management station to receive SNMPv1/v2c traps, and (optionally) event log messages sent as traps from the switch, using the specified community name and destination IPv4 or IPv6 address. You can specify up to ten trap receivers (network management stations). The default community name is **public**.*

[<none | all | non-info | critical | debug>]

(Optional) Configures the security level of the event log messages you want to send as traps to a trap receiver (see table 14-1, “Security Levels for Event Log Messages Sent as Traps”).

- *The type of event log message that you specify applies only to event log messages, not to threshold traps.*
- *For each configured event level, the switch continues to send threshold traps to all network management stations that have the appropriate threshold level configured.*
- *If you do not specify an event level, the switch uses the default value (**none**) and sends no event log messages as traps.*

[<inform>]

(Optional) Configures the switch to send SNMPv2 inform requests when certain events occur. See “Enabling SNMPv2c Informs” on page 14-21 for more information.

Table 14-1. Security Levels for Event Log Messages Sent as Traps

Security Level	Action
None (default)	Sends no event log messages.
All	Sends all event log messages.
Non-Info	Sends all event log messages that are not for information only.
Critical	Sends only event log messages for critical error conditions.
Debug	Sends only event log messages needed to troubleshoot network- and switch-level problems.

For example, to configure a trap receiver in a community named "red-team" with an IP address of 10.28.227.130 to receive only "critical" event log messages, you can enter the following command:

```
ProCurve(config)# snmp-server host 10.28.227.130 red-team  
critical
```

Notes

To replace one community name with another for the same IP address, you must first enter the **no snmp-server host** <community-name> <ipv4-address|ipv6-address> command to delete the unwanted community name. Otherwise, if you add a new community name with an IP address that is already used with a different community name, two valid community name entries are created for the same management station.

If you do not specify the event level ([<none|all|non-info|critical|debug>]), the switch does not send event log messages as traps. However, "well-known" traps and threshold traps (if configured) are still sent.

Enabling SNMPv2c Informs

On a switch enabled for SNMPv2c, you can use the **snmp-server host inform** command to send inform requests when certain events occur. When an SNMP Manager receives an inform request, it can send an SNMP response back to the sending agent on the switch to let the agent know that the inform request reached its destination.

If the sending agent on the switch does not receive an SNMP response back from the SNMP Manager within the timeout period, the inform request may be resent, based on the retry count value.

When you enable SNMPv2c inform requests to be sent, you must specify the IP address and community name of the management station that will receive the inform notification.

Syntax: [no] snmp-server host <ipv4-addr|ipv6-addr> <community name>
inform [retries <count>] [timeout <interval>]]

*Enables (or disables) the **inform** option for SNMPv2c on the switch and allows you to configure options for sending SNMP inform requests.*

retries: *Maximum number of times to resend an inform request if no SNMP response is received. Default: 3*

timeout: *Number of seconds to wait for an acknowledgement before resending the inform request. Default: 15 seconds*

Note

The **retries** and **timeout** values are not used to send trap requests.

To verify the configuration of SNMPv2c informs, enter the **show snmp-server** command:

```
ProCurve Switch 5406zl(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
  ...

Address          Community      Events Sent  Notify Type  Retry  Timeout
-----
15.28.333.456    guest          All          inform        3      15

Excluded MIBs

Snmp Response Pdu Source-IP Information
  Selection Policy : Default rfc1517

Trap Pdu Source-IP Information
  Selection Policy : Configured IP
  Ip Address       : 10.10.10.10
```

SNMPv2c Inform configuration

Figure 14-8. Display of SNMPv2c Inform Configuration

Configuring SNMPv3 Notifications

The SNMPv3 notification process allows messages that are passed via SNMP between the switch and a network management station to be authenticated and encrypted.

To configure SNMPv3 notifications, follow these steps:

1. Enable SNMPv3 operation on the switch by entering the **snmpv3 enable** command (see “SNMP Version 3 Commands” on page 14-6).

When SNMPv3 is enabled, the switch supports:

- Reception of SNMPv3 notification messages (traps and informs)
 - Configuration of initial users
 - (Optional) Restriction of non-SNMPv3 messages to “read only”
2. Configure SNMPv3 users by entering the **snmpv3 user** command (see “SNMPv3 Users” on page 14-7). Each SNMPv3 user configuration is entered in the User Table.
 3. Assign SNMPv3 users to security groups according to their level of access privilege by entering the **snmpv3 group** command (see “Assigning Users to Groups” on page 14-10).
 4. Define the name of an SNMPv3 notification configuration by entering the **snmpv3 notify** command.

Syntax: [no] snmpv3 notify <notify_name> tagvalue <tag_name>

*Associates the name of an SNMPv3 notification configuration with a tag name used (internally) in SNMPv3 commands. To delete a notification-to-tag mapping, enter **no snmpv3 notify <notify_name>**.*

notify <notify_name >

Specifies the name of an SNMPv3 notification configuration.

tagvalue <tag_name >

*Specifies the name of a tag value used in other SNMPv3 commands, such as **snmpv3 targetaddress params taglist <tag_name>** in Step 5.*

5. Configure the target address of the SNMPv3 management station to which SNMPv3 informs and traps are sent by entering the **snmpv3 targetaddress** command.

Syntax: [no] snmpv3 targetaddress < ipv4-addr | ipv6-addr > < name >

Configures the IPv4 or IPv6 address, name, and configuration filename of the SNMPv3 management station to which notification messages are sent.

params < params_name >

*Name of the SNMPv3 station's parameters file. The parameters filename configured with **params** <params_name> must match the **params** <params_name> value entered with the **snmpv3 params** command in Step 6.*

taglist <tag_name> [tag_name] ...

Specifies the SNMPv3 notifications (identified by one or more <tag_name> values) to be sent to the IP address of the SNMPv3 management station.

*You can enter more than one <tag_name> value. Each <tag_name> value must be already associated with the name of an SNMPv3 notification configuration entered with the **snmpv3 notify** command in Step 4.*

Use a blank space to separate <tag_name> values.

*You can enter up to 103 characters in <tag_name> entries following the **taglist** keyword.*

[filter < none | debug | all | not-info | critical >]

*(Optional) Configures the type of messages sent to a management station. Default: **none**.*

[udp-port < port >]

*(Optional) Specifies the UDP port to use. Default: **162**.*

[port-mask < mask >]

*(Optional) Specifies a range of UDP ports. Default: **0**.*

[addr-mask < mask >]

*(Optional) Specifies a range of IP addresses as destinations for notification messages. Default: **0**.*

[retries < value >]

*(Optional) Number of times a notification is retransmitted if no response is received. Range: 1-255. Default: **3**.*

Syntax: [no] snmpv3 targetaddress < ipv4-addr | ipv6-addr > < name >
—Continued—

[timeout < value >]

*(Optional) Time (in millisecond increments) allowed to receive a response from the target before notification packets are retransmitted. Range: 0-2147483647. Default: **1500** (15 seconds).*

[max-msg-size < size >]

*(Optional) Maximum number of bytes supported in a notification message to the specified target. Default: **1472***

6. Create a configuration record for the target address with the **snmpv3 params** command.

Syntax [no] snmpv3 params < params_name > user < user_name >

*Applies the configuration parameters and IP address of an SNMPv3 management station (from the **params** < params_name > value configured with the **snmpv3 targetaddress** command in Step 5) to a specified SNMPv3 user (from the **user** < user_name > value configured with the **snmpv3 user** command in Step 2).*

*If you enter the **snmpv3 params user** command, you must also configure a security model (**sec-model**) and message processing algorithm (**msg-processing**).*

< sec-model < ver1 | ver2c | ver3 >

*Configures the security model used for SNMPv3 notification messages sent to the management station configured with the **snmpv3 targetaddress** command in Step 5.*

*If you configure the security model as **ver3**, you must also configure the message processing value as **ver3**.*

< msg-processing < ver1 | ver2c | ver3 > [noauth | auth | priv]

Configures the algorithm used to process messages sent to the SNMPv3 target address.

*If you configure the message processing value as **ver3** and the security model as **ver3**, you must also configure a security services level (**noauth**, **auth**, or **priv**).*

An example of how to configure SNMPv3 notification is shown here:

The diagram shows a configuration session for a switch. It includes three callout boxes with arrows pointing to specific parts of the configuration commands:

- Top-left callout:** Params_name value in the `snmpv3 targetaddress` command matches the params_name value in the `snmpv3 params` command.
- Top-right callout:** The tag_name value in `snmpv3 notify` command matches the tag_name value in the `snmpv3 targetaddress` command.
- Bottom callout:** Configuring the security model `ver3` requires you to configure message processing `ver3` and a security service level.

```
ProCurve (config)# snmpv3 notify MyNotification tagvalue not_tag
ProCurve (config)# snmpv3 targetaddress not_addr params not_params 15.255.123.109
                    filter not-info taglist not_tag
ProCurve (config)# snmpv3 params not_params user NetworkMgr sec-model ver3
                    message-processing ver3 priv
```

Figure 14-9. Example of an SNMPv3 Notification Configuration

Managing Network Security Notifications

By default, a switch is enabled to send the SNMP notifications listed in “Supported Notifications” on page 14-17 when a network security event (for example, authentication failure) occurs. However, before security notifications can be sent, you must first configure one or more trap receivers or SNMPv3 management stations as described in:

- “Configuring an SNMP Trap Receiver” on page 14-19
- “Configuring SNMPv3 Notifications” on page 14-23

You can manage the default configuration of the switch to disable and re-enable notifications to be sent for the following types of security events:

- ARP protection events
- Unable to establish a connection with the RADIUS or TACACS+ authentication server
- DHCP snooping events
- Dynamic IP Lockdown hardware resources consumed
- Link change notification
- Invalid password entered in a login attempt through a direct serial, Telnet, or SSH connection
- Manager password changes
- Port-security (web, MAC, or 802.1X) authentication failure
- SNMP authentication failure

To enable or disable notification/traps for network security failures and other security events, enter the **snmp-server enable traps** command.

Syntax: [no] snmp-server enable traps [snmp-auth | password-change-mgr | login-failure-mgr | port-security | auth-server-fail | dhcp-snooping | arp-protect]

Enables or disables sending one of the security notification types listed below to configured trap receivers. (Unless otherwise stated, all of the following notifications are enabled in the default configuration.)

- **arp-protect** sends a trap if ARP packets are received with an invalid source or destination MAC address, an invalid IP address, or an invalid IP-to-MAC binding.
- **auth-server-fail** sends a trap if the connection with a RADIUS or TACACS+ authentication server fails.
- **dhcp-snooping** sends a trap if DHCP packets are received from an untrusted source or if DHCP packets contain an invalid IP-to-MAC binding.
- **dyn-ip-lockdown** sends a trap if the switch is out of hardware resources needed to program a dynamic IP lockdown rule.
- **link-change < port-list >** sends a trap when the link state on a port changes from up to down, or the reverse.
- **login-failure-mgr** sends a trap for a failed login with a manager password.
- **password-change-mgr** sends a trap when a manager password is reset.
- **port-security** sends a trap for a failed authentication attempt through a web, MAC, or 801.X authentication session.
- **snmp-authentication [extended | standard]** sends a trap for a failed authentication attempt via SNMP. Default: **extended**.

To determine the specific cause of a security event, check the event log in the console interface to see why a trap was sent. For more information, refer to “Using the Event Log for Troubleshooting Switch Problems” on page C-27.

To display the current configuration for network security notifications, enter the **show snmp-server traps** command. Note that command output is a subset of the information displayed with the **show snmp-server** command in Figure 14-12.

```
ProCurve(config)# show snmp-server traps

Trap Receivers

Link-Change Traps Enabled on Ports [All] : A1-A24

Traps Category                Current Status
-----
SNMP Authentication           : Extended
Password change               : Enabled
Login failures                 : Enabled
Port-Security                 : Enabled
Authorization Server Contact  : Enabled
DHCP Snooping                 : Enabled
Dynamic ARP Protection        : Enabled
Dynamic IP Lockdown           : Enabled

Address            Community  Events Sent  Notify Type  Retry  Timeout
-----
15.255.5.225      public    All          trap          3      15
2001:0db8:0000:0001
:0000:0000:0000:0121 user_1    All          trap          3      15

Excluded MIBs
```

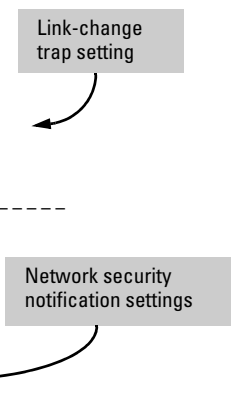


Figure 14-10. Display of Configured Network Security Notifications

Enabling Link-Change Traps

By default a switch is enabled to send a trap when the link state on a port changes from up to down (linkDown) or down to up (linkUp). To reconfigure the switch to send link-change traps to configured trap receivers, enter the **snmp-server enable traps link-change** command.

Syntax: [no] snmp-server enable traps link-change<port-list> [all]

Enables or disables the switch to send a link-change trap to configured trap receivers when the link state on a port goes from up to down or down to up.

*Enter **all** to enable or disable link-change traps on all ports on the switch.*

Configuring the Source IP Address for SNMP Notifications

The switch uses an interface IP address as the source IP address in IP headers when sending SNMP notifications (traps and informs) or responses to SNMP requests.

For multi-netted interfaces, the source IP address is the IP address of the outbound interface of the SNMP reply, which may differ from the destination IP address in the IP header of the received request. For security reasons, it may be desirable to send an SNMP reply with the IP address of the destination interface (or a specified IP address) on which the corresponding SNMP request was received.

To configure the switch to use the source IP address on which an SNMP request was received in SNMP notification/traps and replies, enter the **snmp-server response-source** and **snmp-server trap-source** commands.

Syntax: [no] snmp-server response-source [dst-ip-of-request | <ipv4-addr | ipv6-addr> | loopback<0-7>]

Specifies the source IP address of the SNMP response PDU. The default SNMP response PDU uses the IP address of the active interface from which the SNMP response was sent as the source IP address.

*The **no** form of the command resets the switch to the default behavior (compliant with rfc-1517).*

Default: Interface IP address

dst-ip-of-request: *Destination IP address of the SNMP request PDU that is used as the source IP address in an SNMP response PDU.*

<ipv4-addr | ipv6-addr>: *User-defined interface IP address that is used as the source IP address in an SNMP response PDU. Both IPv4 and IPv6 addresses are supported.*

loopback <0-7>: *IP address configured for the specified loopback interface that is used as the source IP address in an SNMP response PDU. If multiple loopback IP addresses are configured, the lowest alphanumeric address is used.*

For example, to use the IP address of the destination interface on which an SNMP request was received as the source IP address in the IP header of SNMP traps and replies, enter the following command:

```
ProCurve(config)# snmp-server response-source  
dst-ip-of-request
```

To configure the switch to use a specified source IP address in generated trap PDUs, enter the **snmp-server trap-source** command.

Syntax: [no] snmp-server trap-source [<ipv4-addr> | loopback<0-7>]

*Specifies the source IP address to be used for a trap PDU. The **no** form of the command resets the switch to the default behavior (compliant with rfc-1517).*

Default: Use the interface IP address in generated trap PDUs.

<ipv4-addr>: User-defined interface IPv4 address that is used as the source IP address in generated traps. IPv6 addresses are not supported.

***loopback <0-7>**: IP address configured for the specified loopback interface that is used as the source IP address in a generated trap PDU. If multiple loopback IP addresses are configured, the lowest alphanumeric address is used.*

Notes

When you use the **snmp-server response-source** and **snmp-server trap-source** commands, note the following behavior:

- The **snmp-server response-source** and **snmp-server trap-source** commands configure the source IP address for IPv4 interfaces only.
- You must manually configure the **snmp-server response-source** value if you wish to change the default user-defined interface IP address that is used as the source IP address in SNMP traps (RFC 1517).
- The values configured with the **snmp-server response-source** and **snmp-server trap-source** commands are applied globally to all interfaces that are sending SNMP responses or SNMP trap PDUs.
- Only the source IP address field in the IP header of the SNMP response PDU can be changed.
- Only the source IP address field in the IP header and the SNMPv1 Agent Address field of the SNMP trap PDU can be changed.

To verify the configuration of the interface IP address used as the source IP address in IP headers for SNMP replies and traps sent from the switch, enter the **show snmp-server** command to display the SNMP policy configuration.

```
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

...

Excluded MIBs
Snm Response Pdu Source-IP Information
Selection Policy : dstIpOfRequest

Trap Pdu Source-IP Information
Selection Policy : Configured IP
Ip Address       : 10.10.10.10
```

dstIpOfRequest: The destination IP address of the interface on which an SNMP request is received is used as the source IP address in SNMP replies.

Figure 14-11. Display of Source IP Address Configuration

Displaying SNMP Notification Configuration

Use the **show snmp-server** command to display the currently configured:

- Management stations (trap receivers)
- Settings for network security notifications and link-change traps
- SNMP communities

Syntax: show snmp-server

Displays the currently configured notification settings for versions SNMPv1 and SNMPv2c traps, including SNMP communities, trap receivers, link-change traps, and network security notifications.

In the following example, the **show snmp-server** command output shows that the switch has been configured to send SNMP traps and notifications to management stations that belong to the “public”, “red-team”, and “blue-team” communities.

```

ProCurve(config)# show snmp-server

SNMP Communities
Community Name  MIB View Write Access
-----
public          Operator Restricted
blue-team       Manager Unrestricted
red-team        Manager Unrestricted

Trap Receivers

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

Trap Category                Current Trap Configuration
-----
SNMP Authentication          extended
Password change              enabled
Login failures               enabled
Port-Security                enabled
Authorization Server Contact enabled
ARP Protection               enabled
DHCP Snooping                enabled

Address      Community  Events Sent  Notify Type  Retry  Timeout
-----
10.28.227.200 public     All          trap         3      15
10.28.227.105 red-team   Critical     trap         3      15
10.28.227.120 blue-team  Not-INFO     trap         3      15
...
    
```

Figure 14-12. Display of SNMP Notification Configuration

Configuring Listening Mode

For switches that have a separate out-of-band management port, you can specify whether a configured SNMP server listens for SNMP queries over the out-of-band management interface, the data interface, or both. By default, the switch listens over both interfaces.

This option is not available for switches that do not have a separate out-of-band management port. Refer to Appendix I, “Network Out-of-Band Management (OOBM)” in this guide for more information on network out-of-band management.

The listening mode is set with parameters to the **snmp-server** command:

Syntax: `snmp-server [listen <oobm | data | both>]`

Enables or disables inbound SNMP access on a switch.

*Use the **no** version of the command to disable inbound SNMP access.*

*The **listen** parameter is available only on switches that have a separate out-of-band management port. Values for this parameter are:*

- **oobm** — *inbound SNMP access is enabled only on the out-of-band management port.*
- **data** — *inbound SNMP access is enabled only on the data ports.*
- **both** — *inbound SNMP access is enabled on both the out-of-band management port and on the data ports. This is the default value.*

Refer to Appendix I, “Network Out-of-Band Management” in this guide for more information on out-of-band management.

*The **listen** parameter is not available on switches that do not have a separate out-of-band management port.*

Advanced Management: RMON

The switch supports RMON (Remote Monitoring) on all connected network segments. This allows for troubleshooting and optimizing your network.

The following RMON groups are supported:

- Ethernet Statistics (except the numbers of packets of different frame sizes)
- Alarm
- History (of the supported Ethernet statistics)
- Event

The RMON agent automatically runs in the switch. Use the RMON management station on your network to enable or disable specific RMON traps and events. Note that you can access the Ethernet statistics, Alarm, and Event groups from the ProCurve Manager network management software. For more on ProCurve Manager, visit the ProCurve Networking web site at

www.procurve.com

Click on **products index**, then look for the ProCurve Manager topic under the **Network Manager** bar.

CLI-Configured sFlow with Multiple Instances

In earlier software releases, sFlow was configured on the switch via SNMP using a single sFlow instance. Beginning with software release K.11.34, sFlow can also be configured via the CLI for up to three distinct sFlow instances: once enabled, an sFlow receiver/destination can be independently configured for full flow-sampling and counter-polling. CLI-configured sFlow instances may be saved to the startup configuration to persist across a switch reboot.

Terminology

sFlow — An industry standard sampling technology, defined by RFC 3176, used to continuously monitor traffic flows on all ports providing network-wide visibility into the use of the network.

sFlow agent — A software process that runs as part of the network management software within a device. The agent packages data into datagrams that are forwarded to a central data collector.

sFlow destination — The central data collector that gathers datagrams from sFlow-enabled switch ports on the network. The data collector decodes the packet headers and other information to present detailed Layer 2 to Layer 7 usage statistics.

Configuring sFlow

The following sFlow commands allow you to configure sFlow instances via the CLI.

Syntax: [no] sflow <receiver-instance> destination <ip-address> [udp-port-num]

Enables an sFlow receiver/destination. The receiver-instance number must be a 1, 2, or 3. By default, the udp destination port number is 6343.

To disable an sFlow receiver/destination, enter no sflow <receiver-instance>.

Syntax: sflow <receiver-instance> sampling <port-list> <sampling rate>

Once an sFlow receiver/destination has been enabled, this command enables flow sampling for that instance. The receiver-instance number is 1, 2, or 3, and the sampling rate is the allowable non-zero skipcount for the specified port or ports.

To disable flow-sampling for the specified port-list, repeat the above command with a sampling rate of "0".

Syntax: sflow <receiver-instance> polling <port-list> <polling interval>

Once an sFlow receiver/destination has been enabled, this command enables counter polling for that instance. The receiver-instance number is 1, 2, or 3, and the polling interval may be set to an allowable non-zero value to enable polling on the specified port or ports.

To disable counter-polling for the specified port-list, repeat the above command with a polling interval of "0".

Note

Under the multiple instance implementation, sFlow can be configured via the CLI or via SNMP. However, CLI-owned sFlow configurations cannot be modified via SNMP, whereas SNMP-owned instances can be disabled via the CLI using the **no sflow <receiver-instance>** command.

Viewing sFlow Configuration and Status

The following sFlow commands allow you to display sFlow configuration and status via the CLI.

Syntax: show sflow agent

Displays sFlow agent information. The agent address is normally the ip address of the first vlan configured.

Syntax: show sflow <receiver instance> destination

Displays information about the management station to which the sFlow sampling-polling data is sent.

Syntax: show sflow <receiver instance> sampling-polling <port-list/range>

Displays status information about sFlow sampling and polling.

The **show sflow agent** command displays read-only switch agent information. The version information shows the sFlow version, MIB support and software versions; the agent address is typically the ip address of the first vlan configured on the switch.

```
ProCurve# show sflow agent

Version          1.3;HP;K.11.40
Agent Address    10.0.10.228
```

Figure 14-13. Example of Viewing sFlow Agent Information

The **show sflow <instance> destination** command includes information about the management-station's destination address, receiver port, and owner.

```
ProCurve# show sflow 2 destination

Destination Instance      2
sflow                     Enabled
Datagrams Sent           221
Destination Address       10.0.10.41
Receiver Port             6343
Owner                     Administrator, CLI-owned, Instance 2
Timeout (seconds)         99995530
Max Datagram Size         1400
Datagram Version Support  5
```

Figure 14-14. Example of Viewing sFlow Destination Information

Note the following details:

- **Destination Address** remains blank unless it has been configured.
- **Datagrams Sent** shows the number of datagrams sent by the switch agent to the management station since the switch agent was last enabled.
- **Timeout** displays the number of seconds remaining before the switch agent will automatically disable sFlow (this is set by the management station and decrements with time).
- **Max Datagram Size** shows the currently set value (typically a default value, but this can also be set by the management station).

The **show sflow <instance> sampling-polling [port-list]** command displays information about sFlow sampling and polling on the switch. You can specify a list or range of ports for which to view sampling information.

```
ProCurve# show sflow 2 sampling-polling A1-A4
```

Number denotes the sampling/polling instance to which the receiver is coupled.

Port	Sampling			Dropped				Polling							
	Enabled	Rate	Header	Samples				Enabled	Interval						
A1	Yes (2)	40	128	1	2	3	4	5	6	7	8	9	0	---	---
A2	---	---	---	0				Yes (1)	60						
A3	No (1)	0	100	898703				No	30						
A4	Yes (3)	50	128	0				No (3)	0						

Figure 14-15. Example of Viewing sFlow Sampling and Polling Information

Note

The sampling and polling instances (noted in parentheses) coupled to a specific receiver instance are assigned dynamically, and so the instance numbers may not always match. The key thing to note is whether sampling or polling is enabled on a port, and the sampling rates or polling intervals for the receiver instance configured on each port.

LLDP (Link-Layer Discovery Protocol)

To standardize device discovery on all ProCurve switches, LLDP will be implemented while offering limited read-only support for CDP as documented in this manual. For the latest information on your switch model, consult the Release Notes (available on the ProCurve Networking web site). If LLDP has not yet been implemented (or if you are running an older version of software), consult a previous version of the Management and Configuration Guide for device discovery details.

Table 14-2. LLDP and LLDP-MED Features

Feature	Default	Menu	CLI	Web
View the switch's LLDP configuration	n/a	—	page 14-46	—
Enable or disable LLDP on the switch	Enabled	—	page 14-42	—
Change the transmit interval (refresh-interval) for LLDP packets	30 seconds	—	page 14-49	—
Change the holdtime multiplier for LLDP Packets (holdtime-multiplier x refresh-interval = time-to-live)	4 seconds	—	page 14-42	—
Change the delay interval between advertisements	2 seconds	—	page 14-50	—
Changing the reinitialization delay interval	2 seconds	—	page 14-51	—
Configuring SNMP notification support	Disabled	—	page 14-52	—
Configuring transmit and receive modes	tx_rx	—	page 14-53	—
Configuring basic LLDP per-port advertisement content	Enabled	—	page 14-54	—
Configuring port speed and duplex advertisements for optional LLDP and mandatory LLDP-MED applications	Enabled	—	page 14-74	—
Configuring topology change notification for LLDP-MED	Enable	—	page 14-60	—
Changing the fast-start duration for LLDP-MED	5 sec	—	page 14-62	—
Configuring LLDP-MED Advertising	Enabled	—	page 14-54	—
Configuring LLDP-MED device location data	None	—	page 14-72	—
Displaying Advertisement Data and Statistics	n/a	—	page 14-76	—

LLDP (Link Layer Discovery Protocol): provides a standards-based method for enabling the switches covered in this guide to advertise themselves to adjacent devices and to learn about adjacent LLDP devices.

LLDP-MED (LLDP Media Endpoint Discovery): Provides an extension to LLDP and is designed to support VoIP deployments.

Note

LLDP-MED is an extension for LLDP, and the switch requires that LLDP be enabled as a prerequisite to LLDP-MED operation.

An SNMP utility can progressively discover LLDP devices in a network by:

1. Reading a given device's Neighbors table (in the Management Information Base, or MIB) to learn about other, neighboring LLDP devices.
2. Using the information learned in step 1 to find and read the neighbor devices' Neighbors tables to learn about additional devices, and so on.

Also, by using **show** commands to access the switch's neighbor database for information collected by an individual switch, system administrators can learn about other devices connected to the switch, including device type (capability) and some configuration information. In VoIP deployments using LLDP-MED on the switches covered in this guide, additional support unique to VoIP applications is also available. Refer to "LLDP-MED (Media-Endpoint-Discovery)" on page 14-57.

Terminology

Adjacent Device: Refer to "Neighbor or Neighbor Device".

Advertisement: See LLDPDU.

Active Port: A port linked to another active device (regardless of whether MSTP is blocking the link).

ELIN (Emergency Location Identification Number): A valid telephone number in the North American Numbering Plan format and assigned to a multiline telephone system operator by the appropriate authority. This number calls a public service answering point (PSAP) and relays automatic location identification data to the PSAP.

LLDP: Link Layer Discovery Protocol:

- Switches covered in this guide: IEEE 802.1AB

LLDP-Aware: A device that has LLDP in its operating code, regardless of whether LLDP is enabled or disabled.

LLDP Device: A switch, server, router, or other device running LLDP.

LLDP Neighbor: An LLDP device that is either directly connected to another LLDP device or connected to that device by another, non-LLDP Layer 2 device (such as a hub) Note that an 802.1D-compliant switch does not forward LLDP data packets even if it is not LLDP-aware.

LLDPDU (LLDP Data Unit): LLDP data packet are transmitted on active links and include multiple TLVs containing global and per-port switch information. In this guide, LLDPDUs are termed “advertisements” or “packets”.

LLDP-MED (Link Layer Discover Protocol Media Endpoint Discovery): The TIA telecommunications standard produced by engineering subcommittee TR41.4, “VoIP Systems — IP Telephony infrastructure and Endpoints” to address needs related to deploying VoIP equipment in IEEE 802-based environments. This standard will be published as ANSI/TIA-1057.

MIB (Management Information Base): An internal database the switch maintains for configuration and performance information.

MLTS (Multiline Telephone System): A network-based and/or premises-based telephone system having a common interface with the public switched telephone system and having multiple telephone lines, common control units, multiple telephone sets, and control hardware and software.

NANP (North American Numbering Plan): A ten-digit telephone number format where the first three digits are an area code and the last seven-digits are a local telephone number.

Neighbor: See “LLDP Neighbor”.

Non-LLDP Device: A device that is not capable of LLDP operation.

PD (Powered Device): This is an IEEE 802.3af-compliant device that receives its power through a direct connection to a 10/100Base-TX PoE RJ-45 port in a ProCurve fixed-port or chassis-based switch. Examples of PDs include Voice-over-IP (VoIP) telephones, wireless access points, and remote video cameras.

PSAP (Public Safety Answering Point): PSAPs are typically emergency telephone facilities established as a first point to receive emergency (911) calls and to dispatch emergency response services such as police, fire and emergency medical services.

PSE (Power-Sourcing Equipment): A PSE, such as a PoE module installed in a switch covered in this guide, provides power to IEEE 802.3af-compliant PDs directly connected to the ports on the module.

TLV (Type-Length-Value): A data unit that includes a data type field, a data unit length field (in bytes), and a field containing the actual data the unit is designed to carry (as an alphanumeric string, a bitmap, or a subgroup of information). Some TLVs include subelements that occur as separate data points in displays of information maintained by the switch for LLDP advertisements. (That is, some TLVs include multiple data points or subelements.)

General LLDP Operation

An LLDP packet contains data about the transmitting switch and port. The switch advertises itself to adjacent (neighbor) devices by transmitting LLDP data packets out all ports on which outbound LLDP is enabled, and reading LLDP advertisements from neighbor devices on ports that are inbound LLDP-enabled. (LLDP is a one-way protocol and does not include any acknowledgement mechanism.) An LLDP-enabled port receiving LLDP packets inbound from neighbor devices stores the packet data in a Neighbor database (MIB).

LLDP-MED

This capability is an extension to LLDP and is available on the switches covered in this guide. Refer to “LLDP-MED (Media-Endpoint-Discovery)” on page 14-57.

Packet Boundaries in a Network Topology

- Where multiple LLDP devices are directly connected, an outbound LLDP packet travels only to the next LLDP device. An LLDP-capable device does not forward LLDP packets to any other devices, regardless of whether they are LLDP-enabled.
- An intervening hub or repeater forwards the LLDP packets it receives in the same manner as any other multicast packets it receives. Thus, two LLDP switches joined by a hub or repeater handle LLDP traffic in the same way that they would if directly connected.
- Any intervening 802.1D device or Layer-3 device that is either LLDP-unaware or has disabled LLDP operation drops the packet.

Configuration Options

Enable or Disable LLDP on the Switch. In the default configuration, LLDP is globally enabled on the switch. To prevent transmission or receipt of LLDP traffic, you can disable LLDP operation (page 14-42)

Enable or Disable LLDP-MED. In the default configuration for the switches covered in this guide, LLDP-MED is enabled by default. (Requires that LLDP is also enabled.) For more information, refer to “LLDP-MED (Media-Endpoint-Discovery)” on page 14-57.

Change the Frequency of LLDP Packet Transmission to Neighbor Devices. On a global basis, you can increase or decrease the frequency of outbound LLDP advertisements (page 14-42).

Change the Time-To-Live for LLDP Packets Sent to Neighbors. On a global basis, you can increase or decrease the time that the information in an LLDP packet outbound from the switch will be maintained in a neighbor LLDP device (page 14-42).

Transmit and Receive Mode. With LLDP enabled, the switch periodically transmits an LLDP advertisement (packet) out each active port enabled for outbound LLDP transmissions, and receives LLDP advertisements on each active port enabled to receive LLDP traffic (page 14-53). Per-Port configuration options include four modes:

- **Transmit and Receive (tx_rx):** This is the default setting on all ports. It enables a given port to both transmit and receive LLDP packets, and to store the data from received (inbound) LLDP packets in the switch’s MIB.
- **Transmit only (txonly):** This setting enables a port to transmit LLDP packets that can be read by LLDP neighbors. However, the port drops inbound LLDP packets from LLDP neighbors without reading them. This prevents the switch from learning about LLDP neighbors on that port.
- **Receive only (rxonly):** This setting enables a port to receive and read LLDP packets from LLDP neighbors, and to store the packet data in the switch’s MIB. However, the port does not transmit outbound LLDP packets. This prevents LLDP neighbors from learning about the switch through that port.
- **Disable (disable):** This setting disables LLDP packet transmissions and reception on a port. In this state, the switch does not use the port for either learning about LLDP neighbors or informing LLDP neighbors of its presence.

SNMP Notification. You can enable the switch to send a notification to any configured SNMP trap receiver(s) when the switch detects a remote LLDP data change on an LLDP-enabled port (page 14-52).

Per-Port (Outbound) Data Options. The following table lists the information the switch can include in the per-port, outbound LLDP packets it generates. In the default configuration, all outbound LLDP packets include this information in the TLVs transmitted to neighbor devices. However, you can configure LLDP advertisements on a per-port basis to omit some of this information (page 14-54).

Table 14-3. Data Available for Basic LLDP Advertisements

Data Type	Configuration Options	Default	Description
Time-to-Live	See note 1.	120 Seconds	The length of time an LLDP neighbor retains the advertised data before discarding it.
Chassis Type ^{2, 6}	N/A	Always Enabled	Indicates the type of identifier used for Chassis ID.
Chassis ID ⁶	N/A	Always Enabled	Uses base MAC address of the switch.
Port Type ^{3, 6}	N/A	Always Enabled	Uses "Local", meaning assigned locally by LLDP.
Port Id ⁶	N/A	Always Enabled	Uses port number of the physical port. In the switches covered in this guide, this is an internal number reflecting the reserved slot/port position in the chassis. For more information on this numbering scheme, refer to figures D-2 and D-3 in Appendix D, "MAC Address Management" of the <i>Management and Configuration Guide</i> for your switch.
Remote Management Address			
Type ^{4, 6}	N/A	Always Enabled	Shows the network address type.
Address ⁴	Default or Configured	Uses a default address selection method unless an optional address is configured. See "Remote Management Address" on page 14-44.	
System Name ⁶	Enable/Disable	Enabled	Uses the switch's assigned name.
System Description ⁶	Enable/Disable	Enabled	Includes switch model name and running software version, and ROM version.
Port Description ⁶	Enable/Disable	Enabled	Uses the physical port identifier.
System capabilities supported ^{5, 6}	Enable/Disable	Enabled	Identifies the switch's primary capabilities (bridge, router).
System capabilities enabled ^{5, 6}	Enable/Disable	Enabled	Identifies the primary switch functions that are enabled, such as routing.

Data Type	Configuration Options	Default	Description
¹			The Packet Time-to-Live value is included in LLDP data packets. (Refer to “Changing the Time-to-Live for Transmitted Advertisements” on page 14-50.)
²			Subelement of the Chassis ID TLV.
³			Subelement of the Port ID TLV.
⁴			Subelement of the Remote-Management-Address TLV.
⁵			Subelement of the System Capability TLV.
⁶			Populated with data captured internally by the switch. For more on these data types, refer to the IEEE P802.1AB Standard.

Remote Management Address. The switch always includes an IP address in its LLDP advertisements. This can be either an address selected by a default process, or an address configured for inclusion in advertisements. Refer to “IP Address Advertisements” on page 14-45.

Debug Logging. You can enable LLDP debug logging to a configured debug destination (Syslog server and/or a terminal device) by executing the **debug lldp** command. (For more on Debug and Syslog, refer to the “Troubleshooting” appendix in this guide.) Note that the switch’s Event Log does not record usual LLDP update messages.

Options for Reading LLDP Information Collected by the Switch

You can extract LLDP information from the switch to identify adjacent LLDP devices. Options include:

- Using the switch’s **show lldp info** command options to display data collected on adjacent LLDP devices—as well as the local data the switch is transmitting to adjacent LLDP devices (page 14-46).
- Using an SNMP application that is designed to query the Neighbors MIB for LLDP data to use in device discovery and topology mapping. 3400/6400 only?
- Using the **walkmib** command to display a listing of the LLDP MIB objects

LLDP and LLDP-MED Standards Compatibility

The operation covered by this section is compatible with these standards:

- IEEE P802.1AB
- RFC 2922 (PTOPO, or Physical Topology MIB)

- RFC 2737 (Entity MIB)
- RFC 2863 (Interfaces MIB)
- ANSI/TIA-1057/D6 (LLDP-MED; refer to “LLDP-MED (Media-Endpoint-Discovery)” on page 14-57.)

LLDP Operating Rules

(For additional information specific to LLDP-MED operation, refer to “LLDP-MED (Media-Endpoint-Discovery)” on page 14-57.)

Port Trunking. LLDP manages trunked ports individually. That is, trunked ports are configured individually for LLDP operation, in the same manner as non-trunked ports. Also, LLDP sends separate advertisements on each port in a trunk, and not on a per-trunk basis. Similarly, LLDP data received through trunked ports is stored individually, per-port.

IP Address Advertisements. In the default operation, if a port belongs to only one static VLAN, then the port advertises the lowest-order IP address configured on that VLAN. If a port belongs to multiple VLANs, then the port advertises the lowest-order IP address configured on the VLAN with the lowest VID. If the qualifying VLAN does not have an IP address, the port advertises 127.0.0.1 as its IP address. For example, if the port is a member of the default VLAN (VID = 1), and there is an IP address configured for the default VLAN, then the port advertises this IP address. In the default operation, the IP address that LLDP uses can be an address acquired by DHCP or Bootp.

You can override the default operation by configuring the port to advertise any IP address that is manually configured on the switch, even if the port does not belong to the VLAN configured with the selected IP address (page 14-54). (Note that LLDP cannot be configured through the CLI to advertise an addresses acquired through DHCP or Bootp. However, as mentioned above, in the default LLDP configuration, if the lowest-order IP address on the VLAN with the lowest VID for a given port is a DHCP or Bootp address, then the switch includes this address in its LLDP advertisements unless another address is configured for advertisements on that port.) Also, although LLDP allows configuring multiple remote management addresses on a port, only the lowest-order address configured on the port will be included in outbound advertisements. Attempting to use the CLI to configure LLDP with an IP address that is either not configured on a VLAN, or has been acquired by DHCP or Bootp results in the following error message.

```
xxx.xxx.xxx.xxx: This IP address is not configured or is  
a DHCP address.
```

Spanning-Tree Blocking. Spanning tree does not prevent LLDP packet transmission or receipt on STP-blocked links.

802.1X Blocking. Ports blocked by 802.1X operation do not allow transmission or receipt of LLDP packets.

Configuring LLDP Operation

In the default configuration, LLDP is enabled and in both transmit and receive mode on all active ports. The LLDP configuration includes global settings that apply to all active ports on the switch, and per-port settings that affect only the operation of the specified ports.

The commands in this section affect both LLDP and LLDP-MED operation. For information on operation and configuration unique to LLDP-MED, refer to “LLDP-MED (Media-Endpoint-Discovery)” on page 14-57.

Command	Page
show lldp config	14-48
[no] lldp run	14-48
lldp refresh-interval	14-49
lldp holdtime-multiplier	14-50
lldpTxDelay	14-50
lldpReinitDelay	14-51
lldp enable-notification	14-52
lldpnotificationinterval	14-53
lldp admin-status < txonly rxonly tx_rx disable >	14-53
lldp config < port-list > IpAddrEnable	14-54
lldp config < port-list > basicTlvEnable	14-55
lldp config < port-list > dot3TlvEnable < macphy_config >	14-57

Viewing the Current Configuration

Displaying the Global LLDP, Port Admin, and SNMP Notification Status. This command displays the switch’s general LLDP configuration status, including some per-port information affecting advertisement traffic and trap notifications.

Syntax show lldp config

Displays the LLDP global configuration, LLDP port status, and SNMP notification status. For information on port admin status, refer to “Configuring Per-Port Transmit and Receive Modes” on page 14-53.

For example, **show lldp config** produces the following display when the switch is in the default LLDP configuration:

```
ProCurve(config)# show lldp config

LLDP Global Configuration

LLDP Enabled [Yes] : Yes
LLDP Transmit Interval [30] : 30
LLDP Hold time Multiplier [4] : 4
LLDP Delay Interval [2] : 2
LLDP Reinit Interval [2] : 2
LLDP Notification Interval [5] : 5

LLDP Port Configuration

Port | AdminStatus NotificationEnabled
-----+-----
1 | Tx_Rx False
2 | Tx_Rx False
3 | Tx_Rx False
4 | Tx_Rx False
5 | Tx_Rx False
6 | Tx_Rx False
7 | Tx_Rx False
8 | Tx_Rx False
. | .
. | .

Med Topology Trap Enabled
-----
False
True
False
False
True
False
False
```

Note: This value corresponds to the lldp refresh-interval command (page 14-49).

Figure 14-16. Example of Viewing the General LLDP Configuration

Displaying Port Configuration Details. This command displays the port-specific configuration, including.

Syntax show lldp config < port-list >

Displays the LLDP port-specific configuration for all ports in < port-list>, including which optional TLVs and any non-default IP address that are included in the port's outbound advertisements. For information on the notification setting, refer to "Configuring SNMP Notification Support" on page 14-52. For information on the other configurable settings displayed by this command, refer to "Configuring Per-Port Transmit and Receive Modes" on page 14-53.

```
ProCurve(config)# show lldp config a1

LLDP Port Configuration Detail

Port : a1
AdminStatus [Tx_Rx] : Tx_Rx
NotificationEnabled [False] : False
Med Topology Trap Enabled [False] : False

TLVS Advertised:
 * port_descr
 * system_name
 * system_descr
 * system_cap

[ * capabilities
 * network_policy |
 * location_id |
 * poe ]
[ * macphy_config ]

IpAddress Advertised:
```

The diagram shows three callout boxes with arrows pointing to specific parts of the output:

- The first callout box points to the TLV list (* port_descr, * system_name, * system_descr, * system_cap) and contains the text: "These fields appear when medtlvenable is enabled on the switch, which is the default setting."
- The second callout box points to the grouped TLVs (* capabilities, * network_policy, * location_id, * poe, * macphy_config) and contains the text: "This field appears when dot3tlvenable is enabled on the switch, which is the default setting."
- The third callout box points to the "IpAddress Advertised:" line and contains the text: "The blank IpAddress field indicates that the default IP address will be advertised from this port. (Refer to page 14-54: "Configuring a Remote Management Address for Outbound LLDP Advertisements")"

Figure 14-17. Example of Per-Port Configuration Display

Configuring Global LLDP Packet Controls

The commands in this section configure the aspects of LLDP operation that apply the same to all ports in the switch.

Enabling or Disabling LLDP Operation on the Switch. Enabling LLDP operation (the default) causes the switch to:

- Use active, LLDP-enabled ports to transmit LLDP packets describing itself to neighbor devices.

- Add entries to its neighbors table based on data read from incoming LLDP advertisements.

Syntax [no] lldp run

Enables or disables LLDP operation on the switch. The **no** form of the command, regardless of individual LLDP port configurations, prevents the switch from transmitting outbound LLDP advertisements, and causes the switch to drop all LLDP advertisements received from other devices. The switch preserves the current LLDP configuration when LLDP is disabled. After LLDP is disabled, the information in the LLDP neighbors database remains until it times-out. (Default: Enabled)

For example, to disable LLDP on the switch:

```
ProCurve(config)# no lldp run
```

Changing the Packet Transmission Interval. This interval controls how often active ports retransmit advertisements to their neighbors.

Syntax lldp refresh-interval < 5 - 32768 >

Changes the interval between consecutive transmissions of LLDP advertisements on any given port. (Default: 30 seconds)

Note: The **refresh-interval** must be greater than or equal to (4 x **delay-interval**). (The default **delay-interval** is 2). For example, with the default **delay-interval**, the lowest **refresh-interval** you can use is 8 seconds (4 x 2 = 8). Thus, if you want a **refresh-interval** of 5 seconds, you must first change the delay interval to 1 (that is, 4 x 1 < 5). If you want to change the **delay-interval**, use the **setmib** command.

Changing the Time-to-Live for Transmitted Advertisements. The Time-to-Live value (in seconds) for all LLDP advertisements transmitted from a switch is controlled by the switch that generates the advertisement, and determines how long an LLDP neighbor retains the advertised data before discarding it. The Time-to-Live value is the result of multiplying the **refresh-interval** by the **holdtime-multiplier** described below.

Syntax `lldp holdtime-multiplier < 2 - 10 >`

Changes the multiplier an LLDP switch uses to calculate the Time-to-Live for the LLDP advertisements it generates and transmits to LLDP neighbors. When the Time-to-Live for a given advertisement expires the advertised data is deleted from the neighbor switch's MIB. (Default: 4; Range: 2 - 10)

For example, if the refresh-interval on the switch is 15 seconds and the **holdtime-multiplier** is at the default, the Time-to-Live for advertisements transmitted from the switch is 60 seconds (4 x 15). To reduce the Time-to-Live, you could lower the **holdtime-interval** to 2, which would result in a Time-to-Live of 30 seconds.

```
ProCurve(config)# lldp holdtime-multiplier 2
```

Changing the Delay Interval Between Advertisements Generated by Value or Status Changes to the LLDP MIB. The switch uses a *delay-interval* setting to delay transmitting successive advertisements resulting from these LLDP MIB changes. If a switch is subject to frequent changes to its LLDP MIB, lengthening this interval can reduce the frequency of successive advertisements. The delay-interval can be changed using either an SNMP network management application or the CLI **setmib** command.

Syntax `setmib lldpTxDelay.0 -i < 1 - 8192 >`

Uses **setmib** to change the minimum time (delay-interval) any LLDP port will delay advertising successive LLDP advertisements due to a change in LLDP MIB content. (Default: 2; Range: 1 - 8192)

Note: The LLDP refresh-interval (transmit interval) must be greater than or equal to (4 x delay-interval). The switch does not allow increasing the delay interval to a value that conflicts with this relationship. That is, the switch displays **Inconsistent value** if (4 x delay-interval) exceeds the current transmit interval, and the command fails. Depending on the current refresh-interval setting, it may be necessary to increase the refresh-interval before using this command to increase the delay-interval.

For example, to change the delay-interval from 2 seconds to 8 seconds when the refresh-interval is at the default 30 seconds, you must first set the refresh-interval to a minimum of 32 seconds ($32 = 4 \times 8$).

```
ProCurve(config)# setmib lldptxdelay.0 -i 8
lldptxdelay.0: Inconsistent value.
ProCurve(config)# lldp refresh-interval 32
ProCurve(config)# setmib lldptxdelay.0 -i 8
lldpTxDelay.0 = 8
```

Attempt to change the transmit-delay interval shows that the refresh-interval is less than (4 x delay-interval).

Successfully changes the transmit-delay interval to 8.

Changes the refresh-interval to 32; that is: $32 = 4 \times (\text{desired transmit-delay interval})$

Figure 14-18. Example of Changing the Transmit-Delay Interval

Changing the Reinitialization Delay Interval. In the default configuration, a port receiving a **disable** command followed immediately by a **txonly**, **rxonly**, or **tx_rx** command delays reinitializing for two seconds, during which time LLDP operation remains disabled. If an active port is subjected to frequent toggling between the LLDP disabled and enabled states, LLDP advertisements are more frequently transmitted to the neighbor device. Also, the neighbor table in the adjacent device will change more frequently, as it deletes, then replaces LLDP data for the affected port which, in turn, generates SNMP traps (if trap receivers and SNMP notification are configured). All of this can unnecessarily increase network traffic. Extending the reinitialization-

delay interval delays the port's ability to reinitialize and generate LLDP traffic following an LLDP disable/enable cycle.

Syntax `setmib lldpReinitDelay.0 -i < 1 - 10 >`

Uses **setmib** to change the minimum time (reinitialization delay interval) an LLDP port will wait before reinitializing after receiving an LLDP disable command followed closely by a `txonly` or `tx_rx` command. The delay interval commences with execution of the **lldp admin-status < port-list > disable** command. (Default: 2 seconds; Range: 1 - 10 seconds)

For example, the following command changes the reinitialization delay interval to five seconds:

```
ProCurve(config)# setmib lldpreinitdelay.0 -i 5
```

Configuring SNMP Notification Support

You can enable SNMP trap notification of LLDP data changes detected on advertisements received from neighbor devices, and control the interval between successive notifications of data changes on the same neighbor.

Enabling LLDP Data Change Notification for SNMP Trap Receivers.

Syntax `[no] lldp enable-notification < port-list >`

Enables or disables each port in `< port-list >` for sending notification to configured SNMP trap receiver(s) if an LLDP data change is detected in an advertisement received on the port from an LLDP neighbor. (Default: Disabled)

For information on configuring trap receivers in the switch, refer to "SNMP Notifications" on page 14-17.

For example, this command enables SNMP notification on ports 1 - 5:

```
ProCurve(config)# lldp enable-notification 1-5
```

Changing the Minimum Interval for Successive Data Change Notifications for the Same Neighbor.

If LLDP trap notification is enabled on a port, a rapid succession of changes in LLDP information received in advertisements from one or more neighbors can generate a high number of traps. To reduce this effect, you can globally change the interval between successive notifications of neighbor data change.

Syntax `setmib lldpnotificationinterval.0 -i < 1 - 3600 >`

Globally changes the interval between successive traps generated by the switch. If multiple traps are generated in the specified interval, only the first trap will be sent. The remaining traps will be suppressed. (A network management application can periodically check the switch MIB to detect any missed change notification traps. Refer to IEEE P802.1AB or later for more information.) (Default: 5 seconds)

For example, the following command limits change notification traps from a particular switch to one per minute.

```
ProCurve(config)# setmib lldpnotificationinterval.0 -i 60  
lldpNotificationInterval.0 = 60
```

Configuring Per-Port Transmit and Receive Modes

These commands control advertisement traffic inbound and outbound on active ports.

Syntax `lldp admin-status < port-list > < txonly | rxonly | tx_rx | disable >`

With LLDP enabled on the switch in the default configuration, each port is configured to transmit and receive LLDP packets. These options enable you to control which ports participate in LLDP traffic and whether the participating ports allow LLDP traffic in only one direction or in both directions.

txonly: *Configures the specified port(s) to transmit LLDP packets, but block inbound LLDP packets from neighbor devices.*

rxonly: *Configures the specified port(s) to receive LLDP packets from neighbors, but block outbound packets to neighbors.*

tx_rx: *Configures the specified port(s) to both transmit and receive LLDP packets. (This is the default setting.)*

disable: *Disables LLDP packet transmit and receive on the specified port(s).*

Configuring Basic LLDP Per-Port Advertisement Content

In the default LLDP configuration, outbound advertisements from each port on the switch include both mandatory and optional data.

Mandatory Data. An active LLDP port on the switch always includes the mandatory data in its outbound advertisements. LLDP collects the mandatory data, and, except for the Remote Management Address, you cannot use LLDP commands to configure the actual data.

- Chassis Type (TLV subelement)
- Chassis ID (TLV)
- Port Type (TLV subelement)
- Port ID (TLV)
- Remote Management Address (TLV; actual IP address is a subelement that can be a default address or a configured address)

Configuring a Remote Management Address for Outbound LLDP Advertisements. This is an optional command you can use to include a specific IP address in the outbound LLDP advertisements for specific ports.

Syntax [no] lldp config < port-list > ipAddrEnable < ip-address >

Replaces the default IP address for the port with an IP address you specify. This can be any IP address configured in a static VLAN on the switch, even if the port does not belong to the VLAN configured with the selected IP address. The **no** form of the command deletes the specified IP address. If there are no IP addresses configured as management addresses, then the IP address selection method returns to the default operation. (Default: The port advertises the IP address of the lowest-numbered VLAN (VID) to which it belongs. If there is no IP address configured on the VLAN(s) to which the port belongs, and the port is not configured to advertise an IP address from any other (static) VLAN on the switch, then the port advertises an address of 127.0.0.1.)

Note: This command does not accept either IP addresses acquired through DHCP or Bootp, or IP addresses that are not configured in a static VLAN on the switch

For example, if port 3 belongs to a subnetted VLAN that includes an IP address of 10.10.10.100 and you wanted port 3 to use this secondary address in LLDP advertisements, you would need to execute the following command:

```
ProCurve(config)# lldp config 3 ipAddrEnable 10.10.10.100
```

Optional Data. You can configure an individual port or group of ports to exclude one or more of these data types from outbound LLDP advertisements. Note that optional data types, when enabled, are populated with data internal to the switch; that is, you cannot use LLDP commands to configure their actual content.

- port description (TLV)
- system name (TLV)
- system description (TLV)
- system capabilities (TLV)
 - system capabilities Supported (TLV subelement)
 - system capabilities Enabled (TLV subelement)
- port speed and duplex (TLV subelement)

Syntax: [no] lldp config < port-list > basicTlvEnable < TLV-Type >

port_descr

For outbound LLDP advertisements, this TLV includes an alphanumeric string describing the port.
(Default: Enabled)

system_name

For outbound LLDP advertisements, this TLV includes an alphanumeric string showing the system's assigned name.
(Default: Enabled)

system_descr

For outbound LLDP advertisements, this TLV includes an alphanumeric string describing the full name and version identification for the system's hardware type, software version, and networking application.
(Default: Enabled)

system_cap

*For outbound advertisements, this TLV includes a bitmask of supported system capabilities (device functions). Also includes information on whether the capabilities are enabled.
(Default: Enabled)*

For example, if you wanted to exclude the system name TLV from the outbound LLDP advertisements for all ports on a switch, you would use this command:

```
ProCurve(config)# no lldp config 1-24 basicTlvEnable  
system_name
```

If you later decided to reinstate the system name TLV on ports 1-5, you would use this command:

```
ProCurve(config)# lldp config 1-5 basicTlvEnable  
system_name
```

Configuring Support for Port Speed and Duplex Advertisements

This feature is optional for LLDP operation, but is *required* for LLDP-MED operation.

Port speed and duplex advertisements are supported on the switches covered in this guide to inform an LLDP endpoint and the switch port of each other's port speed and duplex configuration and capabilities. Configuration mismatches between a switch port and an LLDP endpoint can result in excessive collisions and voice quality degradation. LLDP enables discovery of such mismatches by supporting SNMP access to the switch MIB for comparing the current switch port and endpoint settings. (Changing a current device configuration to eliminate a mismatch requires intervention by the system operator.)

Syntax: [no] lldp config < port-list > dot3TlvEnable macphy_config

For outbound advertisements, this TLV includes the (local) switch port's current speed and duplex settings, the range of speed and duplex settings the port supports, and the method required for reconfiguring the speed and duplex settings on the device (auto-negotiation during link initialization, or manual configuration).

*Using SNMP to compare local and remote information can help in locating configuration mismatches.
(Default: Enabled)*

Note: For LLDP operation, this TLV is optional. For LLDP-MED operation, this TLV is mandatory.

As mentioned above, an SNMP network management application can be used to compare the port speed and duplex data configured in the switch and advertised by the LLDP endpoint. You can also use the CLI to display this information. For more on using the CLI to display port speed and duplex information, refer to “Displaying the Current Port Speed and Duplex Configuration on a Switch Port” on page 14-73.

LLDP-MED (Media-Endpoint-Discovery)

LLDP-MED (ANSI/TIA-1057/D6) extends the LLDP (IEEE 802.1AB) industry standard to support advanced features on the network edge for Voice Over IP (VoIP) endpoint devices with specialized capabilities and LLDP-MED standards-based functionality. LLDP-MED in the switches uses the standard LLDP commands described earlier in this section, with some extensions, and also introduces new commands unique to LLDP-MED operation. The **show** commands described elsewhere in this section are applicable to both LLDP and LLDP-MED operation. LLDP-MED benefits include:

- plug-and-play provisioning for MED-capable, VoIP endpoint devices
- simplified, vendor-independent management enabling different IP telephony systems to interoperate on one network
- automatic deployment of convergence network policies (voice VLANs, Layer 2/CoS priority, and Layer 3/QoS priority)
- configurable endpoint location data to support the Emergency Call Service (ECS) (such as Enhanced 911 service, 999, 112)
- detailed VoIP endpoint data inventory readable via SNMP from the switch

Configuring for Network Management Applications

LLDP (Link-Layer Discovery Protocol)

- Power over Ethernet (PoE) status and troubleshooting support via SNMP
- support for IP telephony network troubleshooting of call quality issues via SNMP

This section describes how to configure and use LLDP-MED features in the switches to support VoIP network edge devices (Media Endpoint Devices) such as:

- IP phones
- voice/media gateways
- media servers
- IP communications controllers
- other VoIP devices or servers

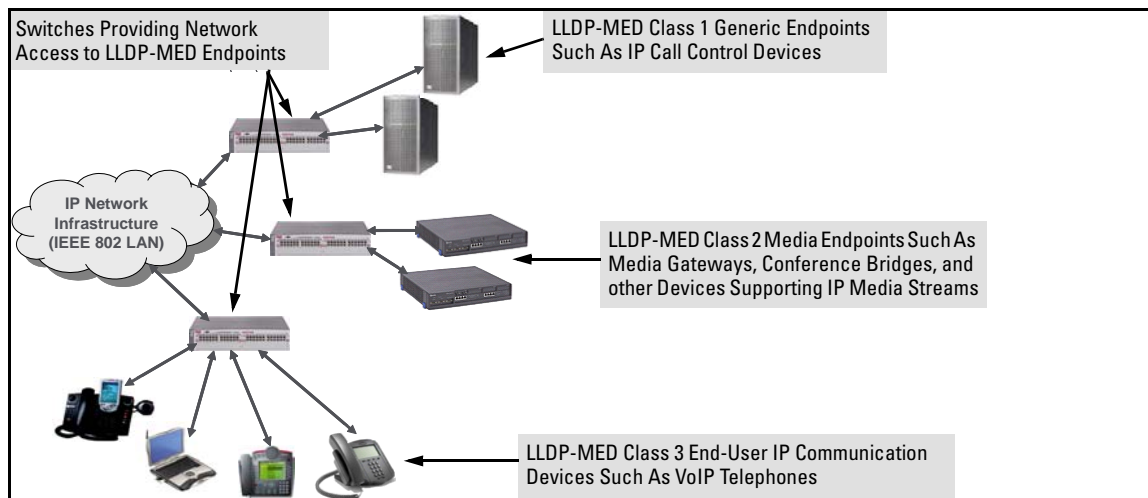


Figure 14-19. Example of LLDP-MED Network Elements

LLDP-MED Endpoint Support. LLDP-MED on the switches covered in this guide interoperates with directly connected IP telephony (endpoint) clients having these features and services:

- able to autonegotiate speed and duplex configuration with the switch

- able to use the following network policy elements configured on the client port
 - voice VLAN ID
 - 802.1p (Layer 2) QoS
 - Diffserv codepoint (DSCP) (Layer 3) QoS
- discover and advertise device location data learned from the switch
- support emergency call service (ECS—such as E911, 999, and 112)
- advertise device information for the device data inventory collected by the switch, including:
 - hardware revision
 - serial number
 - asset ID
 - firmware revision
 - manufacturer name
 - software revision
 - model name
- provide information on network connectivity capabilities (for example, a multi-port VoIP phone with Layer 2 switch capability)
- support the fast start capability

Note

LLDP-MED on the switches covered in this guide is intended for use with VoIP endpoints, and is not designed to support links between network infrastructure devices, such as switch-to-switch or switch-to-router links.

LLDP-MED Endpoint Device Classes. LLDP-MED endpoint devices are, by definition, located at the network edge and communicate using the LLDP-MED framework. Any LLDP-MED endpoint device belongs to one of the following three classes:

- Class 1 (Generic Endpoint Devices): These devices offer the basic LLDP discovery services, network policy advertisement (VLAN ID, Layer 2/802.1p priority, and Layer 3/DSCP priority), and PoE management. This class includes such devices as IP call controllers and communication-related servers.
- Class 2 (Media Endpoint Devices): These devices offer all Class 1 features plus media streaming capability, and include such devices as voice/media gateways, conference bridges, and media servers.

Configuring for Network Management Applications

LLDP (Link-Layer Discovery Protocol)

- **Class 3 (Communication Devices):** These devices are typically IP phones or end-user devices that otherwise support IP media and offer all Class 1 and Class 2 features, plus location identification and emergency 911 capability, Layer 2 switch support, and device information management.

LLDP-MED Operational Support. The switches covered in this guide offer two configurable TLVs supporting MED-specific capabilities:

- `medTlvEnable` (for per-port enabling or disabling of LLDP-MED operation)
- `medPortLocation` (for configuring per-port location or emergency call data)

Note

LLDP-MED operation also requires the port speed and duplex TLV (`dot3TlvEnable`; page 14-57), which is enabled in the default configuration.

LLDP-MED Topology Change Notification

This optional feature provides information an SNMP application can use to track LLDP-MED connects and disconnects.

Syntax: `lldp top-change-notify < port-list >`

Topology change notification, when enabled on an LLDP port, causes the switch to send an SNMP trap if it detects LLDP-MED endpoint connection or disconnection activity on the port, or an age-out of the LLDP-MED neighbor on the port. The trap includes the following information:

- *the port number (internal) on which the activity was detected (For more in internal port numbers, refer to “Determining the Switch Port Number Included in Topology Change Notification Traps” on page 14-79.)*
- *the LLDP-MED class of the device detected on the port (“LLDP-MED Endpoint Device Classes” on page 14-59.)*

The **show running** command shows whether the topology change notification feature is enabled or disabled. For example, if ports A1-A10 have topology change notification enabled, the following entry appears in the **show running** output:

```
lldp top-change-notify A1-A10
```

(Default: Disabled)

Note: To send traps, this feature requires access to at least one SNMP server. For information on configuring traps, refer to “SNMP Notifications” on page 14-17.

Also, if a detected LLDP-MED neighbor begins sending advertisements without LLDP-MED TLVs, the switch sends a top-change-notify trap.

Note

Topology change notifications provide one method for monitoring system activity. However, because SNMP normally employs UDP, which does not guarantee datagram delivery, topology change notification should not be relied upon as the sole method for monitoring critical endpoint device connectivity.

LLDP-MED Fast Start Control

Syntax: `lldp fast-start-count < 1 - 10 >`

*An LLDP-MED device connecting to a switch port may use the data contained in the MED TLVs from the switch to configure itself. However, the **lldp refresh-interval** setting (default: 30 seconds) for transmitting advertisements can cause an unacceptable delay in MED device configuration. To support rapid LLDP-MED device configuration, the **lldp fast-start-count** command temporarily overrides the **refresh-interval** setting for the **fast-start-count** advertisement interval. This results in the port initially advertising LLDP-MED at a faster rate for a limited time. Thus, when the switch detects a new LLDP-MED device on a port, it transmits one LLDP-MED advertisement per second out the port for the duration of the **fast-start-count** interval. In most cases, the default setting should provide an adequate **fast-start-count** interval.*

(Range: 1 - 10 seconds; Default: 5 seconds)

Note: This global command applies only to ports on which a new LLDP-MED device is detected. It does not override the refresh-interval setting on ports where non-MED devices are detected.

Advertising Device Capability, Network Policy, PoE Status and Location Data

The `medTlvEnable` option on the switch is enabled in the default configuration and supports the following LLDP-MED TLVs:

- LLDP-MED capabilities: This TLV enables the switch to determine:
 - whether a connected endpoint device supports LLDP-MED
 - which specific LLDP-MED TLVs the endpoint supports
 - the device class (1, 2, or 3) for the connected endpoint

This TLV also enables an LLDP-MED endpoint to discover what LLDP-MED TLVs the switch port currently supports.

- network policy operating on the port to which the endpoint is connected (VLAN, Layer 2 QoS, Layer 3 QoS)
- PoE (MED Power-over-Ethernet)
- physical location data — page 14-66

Note

LLDP-MED operation requires the `macphy_config` TLV subelement—enabled by default—that is optional for IEEE 802.1AB LLDP operation. Refer to the `dot3TlvEnable macphy_config` command on page 14-57.

Network Policy Advertisements. Network policy advertisements are intended for real-time voice and video applications, and include these TLV subelements:

- Layer 2 (802.1p) QoS
- Layer 3 DSCP (diffserv code point) QoS
- Voice VLAN ID (VID)

VLAN Operating Rules. These rules affect advertisements of VLANs in network policy TLVs:

- The VLAN ID TLV subelement applies only to a VLAN configured for voice operation (`vlan < vid > voice`).
- If there are multiple voice VLANs configured on a port, LLDP-MED advertises the voice VLAN having the lowest VID.
- The voice VLAN port membership configured on the switch can be tagged or untagged. However, if the LLDP-MED endpoint expects a tagged membership when the switch port is configured for untagged, or the reverse, then a configuration mismatch results. (Typically, the endpoint expects the switch port to have a tagged voice VLAN membership.)
- If a given port does not belong to a voice VLAN, then the switch does not advertise the VLAN ID TLV through this port.

Policy Elements. These policy elements may be statically configured on the switch or dynamically imposed during an authenticated session on the switch using a RADIUS server and 802.1X or MAC authentication. (Web authentication does not apply to VoIP telephones and other telecommunications devices that are not capable of accessing the switch through a Web browser.) The QoS and voice VLAN policy elements can be statically configured with the following CLI commands:

```
vlan < vid > voice
vlan < vid > < tagged | untagged > < port-list >
int < port-list > qos priority < 0 - 7 >
vlan < vid > qos dscp < codepoint >
```

Notes

A codepoint must have an 802.1p priority before you can configure it for use in prioritizing packets by VLAN-ID. If a codepoint you want to use shows **No Override** in the **Priority** column of the DSCP policy table (display with **show qos-dscp map**, then use **qos-dscp map < codepoint > priority < 0 - 7 >** to configure a priority before proceeding. For more on this topic, refer to the chapter titled “Quality of Service (QoS): Managing Bandwidth More Effectively” in the *Advanced Traffic Management Guide* for your switch.

Enabling or Disabling medTlvEnable. In the default LLDP-MED configuration, the TLVs controlled by medTlvEnable are enabled.

Syntax: [no] lldp config < port-list > medTlvEnable < medTlv >

- *Enables or disables advertisement of the following TLVs on the specified ports:*
 - *device capability TLV*
 - *configured network policy TLV*
 - *configured location data TLV (Refer to “Configuring Location Data for LLDP-MED Devices” on page 14-66.)*
 - *current PoE status TLV*

(Default: All of the above TLVs are enabled.)

- *Helps to locate configuration mismatches by allowing use of an SNMP application to compare the LLDP-MED configuration on a port with the LLDP-MED TLVs advertised by a neighbor connected to that port.*

capabilities

This TLV enables the switch to determine:

- *which LLDP-MED TLVs a connected endpoint can discover*
- *the device class (1, 2, or 3) for the connected endpoint*

This TLV also enables an LLDP-MED endpoint to discover what LLDP-MED TLVs the switch port currently supports.

(Default: enabled)

Note: This TLV cannot be disabled unless the network_policy, poe, and location_id TLVs are already disabled.

network-policy

This TLV enables the switch port to advertise its configured network policies (voice VLAN, Layer 2 QoS, Layer 3 QoS), and allows LLDP-MED endpoint devices to auto-configure the voice network policy advertised by the switch. This also enables the use of SNMP applications to troubleshoot statically configured endpoint network policy mismatches.

(Default: Enabled)

Notes: *Network policy is only advertised for ports that are configured as members of the voice VLAN. If the port belongs to more than one voice VLAN, then the voice VLAN with the lowest-numbered VID is selected as the VLAN for voice traffic. Also, this TLV cannot be enabled unless the capability TLV is already enabled.*

For more information, refer to “Network Policy Advertisements” on page 14-63

location_id

This TLV enables the switch port to advertise its configured location data (if any). For more on configuring location data, refer to “Configuring Location Data for LLDP-MED Devices”.

(Default: Enabled)

Note: *When disabled, this TLV cannot be enabled unless the capability TLV is already enabled.*

poe

This TLV enables the switch port to advertise its current PoE (Power over Ethernet) state and to read the PoE requirements advertised by the LLDP-MED endpoint device connected to the port.

(Default: Enabled)

Note: *When disabled, this TLV cannot be enabled unless the capability TLV is already enabled.*

For more on this topic, refer to “PoE Advertisements”, below.

PoE Advertisements. These advertisements inform an LLDP-MED endpoint of the power (PoE) configuration on switch ports. Similar advertisements from an LLDP-MED endpoint inform the switch of the endpoint's power needs and provide information that can be used to identify power priority mismatches.

Power-over-Ethernet TLVs include the following power data:

- **power type:** indicates whether the device is a power-sourcing entity (PSE) or a powered device (PD). Ports on the J8702A PoE zl module are PSE devices. A MED-capable VoIP telephone is a PD.
- **power source:** indicates the source of power in use by the device. Power sources for powered devices (PDs) include PSE, local (internal), and PSE/local. The switches covered in this guide advertise Unknown.
- **power priority:** indicates the power priority configured on the switch (PSE) port or the power priority configured on the MED-capable endpoint.
- **power value:** indicates the total power in watts that a switch port (PSE) can deliver at a particular time, or the total power in watts that the MED endpoint (PD) requires to operate.

To display the current power data for an LLDP-MED device connected to a port, use the following command:

```
show lldp info remote-device < port-list >
```

For more on this command, refer to page 14-74.

To display the current PoE configuration on the switch, use the following commands:

```
show power brief < port-list >
```

```
show power < port-list >
```

For more on PoE configuration and operation, refer to Chapter 11, “Power Over Ethernet (PoE/PoE+) Operation”.

Configuring Location Data for LLDP-MED Devices

You can configure a switch port to advertise location data for the switch itself, the physical wall-jack location of the endpoint (recommended), or the location of a DHCP server supporting the switch and/or endpoint. You also have the option of configuring these different address types:

- **civic address:** physical address data such as city, street number, and building information

- **ELIN (Emergency Location Identification Number):** an emergency number typically assigned to MLTS (Multiline Telephone System Operators) in North America
- **coordinate-based location:** attitude, longitude, and altitude information (Requires configuration via an SNMP application.)

Syntax: [no] lldp config < port-list > medPortLocation < Address-Type >

*Configures location or emergency call data the switch advertises per port in the **location_id** TLV. This TLV is for use by LLDP-MED endpoints employing location-based applications.*

Note: *The switch allows one medPortLocation entry per port (without regard to type). Configuring a new medPortLocation entry of any type on a port replaces any previously configured entry on that port.*

civic-addr < COUNTRY-STR > < WHAT > < CA-TYPE > < CA-VALUE > ...
[< CA-TYPE > < CA-VALUE >] ... [< CA-TYPE > < CA-VALUE >]

This command enables configuration of a physical address on a switch port, and allows up to 75 characters of address information.

COUNTRY-STR: *A two-character country code, as defined by ISO 3166. Some examples include **FR** (France), **DE** (Germany), and **IN** (India). This field is required in a **civic-addr** command. (For a complete list of country codes, visit www.iso.org on the world wide web.)*

WHAT: *A single-digit number specifying the type of device to which the location data applies:*

0: *Location of DHCP server*

1: *Location of switch*

2: *Location of LLDP-MED endpoint (recommended application)*

*This field is required in a **civic-addr** command.*

—Continued—

— Continued—

Type/Value Pairs (CA-TYPE and CA-VALUE): This is a series of data pairs, each composed of a location data “type” specifier and the corresponding location data for that type. That is, the first value in a pair is expected to be the civic address “type” number (**CA-TYPE**), and the second value in a pair is expected to be the corresponding civic address data (**CA-VALUE**). For example, if the **CA-TYPE** for “city name” is “3”, then the type/value pair to define the city of Paris is “**3 Paris**”. Multiple type/value pairs can be entered in any order, although it is recommended that multiple pairs be entered in ascending order of the **CA-TYPE**. When an emergency call is placed from a properly configured class 3 endpoint device to an appropriate PSAP, the country code, device type, and type/value pairs configured on the switch port are included in the transmission. The “type” specifiers are used by the PSAP to identify and organize the location data components in an understandable format for response personnel to interpret. A **civic-addr** command requires a minimum of one type/value pair, but typically includes multiple type/value pairs as needed to configure a complete set of data describing a given location. **CA-TYPE:** This is the first entry in a type/value pair, and is a number defining the type of data contained in the second entry in the type/value pair (**CA-VALUE**). Some examples of **CA-TYPE** specifiers include:

- 3 = city
- 6 = street (name)
- 25 = building name

(Range: 0 - 255)

For a sample listing of **CA-TYPE** specifiers, refer to table 14-4 on page 14-70.

CA-VALUE: This is the second entry in a type/value pair, and is an alphanumeric string containing the location information corresponding to the immediately preceding **CA-TYPE** entry. Strings are delimited by either blank spaces, single quotes (‘...’), or double quotes (“...”). Each string should represent a specific data type in a set of unique type/value pairs comprising the description of a location, and each string must be preceded by a **CA-TYPE** number identifying the type of data in the string.

Note: A switch port allows one instance of any given **CA-TYPE**. For example, if a type/value pair of **6 Atlantic** (to specify “Atlantic” as a street name) is configured on port A5 and later another type/value pair of **6 Pacific** is configured on the same port, then **Pacific** replaces **Atlantic** in the civic address location configured for port A5.

elin-addr < emergency-number >

This feature is intended for use in Emergency Call Service (ECS) applications to support class 3 LLDP-MED VoIP telephones connected to a switch covered in this guide in a multiline telephone system (MLTS) infrastructure. An ELIN (Emergency Location Identification Number) is a valid North American Numbering Plan (NANP) format telephone number assigned to MLTS operators in North America by the appropriate authority. The ELIN is used to route emergency (E911) calls to a Public Safety Answering Point (PSAP).

(Range: 1-15 numeric characters)

Configuring Coordinate-Based Locations. Latitude, longitude, and altitude data can be configured per switch port using an SNMP management application. For more information, refer to the documentation provided with the application. A further source of information on this topic is *RFC 3825-Dynamic Host Configuration Protocol Option for Coordinate-based Location Configuration Information*.

Note

Endpoint use of data from a medPortLocation TLV sent by the switch is device-dependent. Refer to the documentation provided with the endpoint device.

Table 14-4. Some Location Codes Used in CA-TYPE Fields*

Location Element	Code	Location Element	Code
national subdivision	1	street number	19
regional subdivision	2	additional location data	22
city or township	3	unit or apartment	26
city subdivision	4	floor	27
street	6	room number	28
street suffix	18		

*The code assignments in this table are examples from a work-in-progress (the internet draft titled "Dynamic Host Configuration Protocol (DHCPv4 and DHCPv6) Option for Civic Addresses Configuration Information draft-ietf-geopriv-dhcp-civil-06" dated May 30, 2005.) For the actual codes to use, contact the PSAP or other authority responsible for specifying the civic addressing data standard for your network.

Example of a Location Configuration. Suppose a system operator wanted to configure the following information as the civic address for a telephone connected to her company's network through port A2 of a switch at the following location:

Description	CA-Type	CA-VALUE
national subdivision	1	CA
city	3	Widgitville
street	6	Main
street number	19	1433
unit	26	Suite 4-N
floor	27	4
room number	28	N4-3

Figure 14-20 shows the commands for configuring and displaying the above data.

```

ProCurve(config)# lldp config a2 medportlocation civic-addr US 2 1 CA
lle 6 Main 19 1433 26 Suite_4-N 27 4 28 N4-3
ProCurve(config)# show lldp config a2

LLDP Port Configuration Detail

Port : A2
AdminStatus [Tx_Rx] : Tx_Rx
NotificationEnabled [False] : False
Med Topology Trap Enabled [False] : False

Country Name      : US
What              : 2
Ca-Type           : 1
Ca-Length         : 2
Ca-Value          : CA
Ca-Type           : 3
Ca-Length         : 11
Ca-Value          : Widgitville
Ca-Type           : 6
Ca-Length         : 4
Ca-Value          : Main
Ca-Type           : 19
Ca-Length         : 4
Ca-Value          : 1433
Ca-Type           : 26
Ca-Length         : 9
Ca-Value          : Suite_4-N
Ca-Type           : 27
Ca-Length         : 1
Ca-Value          : 4
Ca-Type           : 28
  
```

Figure 14-20. Example of a Civic Address Configuration

Displaying Advertisement Data

Command	Page
show lldp info local-device	below
walkmib lldpXdot3LocPortOperMauType	
show lldp info remote-device	14-74
walkmib lldpXdot3RemPortAutoNegAdvertisedCap	
show lldp info stats	14-76

Displaying Switch Information Available for Outbound Advertisements

These commands display the current switch information that will be used to populate outbound LLDP advertisements.

Syntax `show lldp info local-device [port-list]`

Without the [port-list] option, this command displays the global switch information and the per-port information currently available for populating outbound LLDP advertisements.

With the [port-list] option, this command displays only the following port-specific information that is currently available for outbound LLDP advertisements on the specified ports:

- **PortType**
- **PortId**
- **PortDesc**

Note: This command displays the information available on the switch. Use the **lldp config < port-list >** command to change the selection of information that is included in actual outbound advertisements. In the default LLDP configuration, all information displayed by this command is transmitted in outbound advertisements.

For example, in the default configuration, the switch information currently available for outbound LLDP advertisements appears similar to the display in Figure 14-21 on page 14-73.


```

ProCurve(config)# show lldp info local-device

LLDP Local Device Information

Chassis Type : mac-address
Chassis Id   : 00 08 83 08 db 20
System Name  : ProCurve
System Description : HP J8697A ProCurve Switch 5406zl revision K.11.00 RO...
System Capabilities Supported : bridge, router
System Capabilities Enabled   : bridge
Management Address :
| Type:ipv4          |
|_ Address: _ _ _ _ |

LLDP Port Information

Port      | PortType  PortId  PortDesc
-----+-----+-----+-----
1         | local     1       1
2         | local     2       2
3         | local     3       3
4         | local     4       4
5         | local     5       5
6         | local     6       6
.         | .         .       .
.         | .         .       .
.         | .         .       .

```

The Management Address field displays only the LLDP-configurable IP addresses on the switch. (Only manually-configured IP addresses are LLDP-configurable.) If the switch has only an IP address from a DHCP or Bootp server, then the Management Address field is empty (because there are no LLDP-configurable IP addresses available). For more on this topic, refer to "Remote Management Address" on page 14-44.

Figure 14-21. Example of Displaying the Global and Per-Port Information Available for Outbound Advertisements

```

ProCurve (config)# show lldp info local 1-2

LLDP Local Port Information Detail

Port      : 1
PortType  : local
PortId    : 1
PortDesc  : 1

-----

Port      : 2
PortType  : local
PortId    : 2
PortDesc  : 2

```

Figure 14-22. Example of the Default Per-Port Information Content for Ports 1 and 2

Displaying the Current Port Speed and Duplex Configuration on a Switch Port. Port speed and duplex information for a switch port and a connected LLDP-MED endpoint can be compared for configuration mismatches by using an SNMP application. You can also use the switch CLI to display this information, if necessary. The following two commands provide methods for displaying speed and duplex information for switch ports. For

information on displaying the currently configured port speed and duplex on an LLDP-MED endpoint, refer to “Displaying the Current Port Speed and Duplex Configuration on a Switch Port” on page 14-73.

Syntax: show interfaces brief < port-list >

*Includes port speed and duplex configuration in the **Mode** column of the resulting display.*

Displaying Advertisements Currently in the Neighbors MIB. These commands display the content of the inbound LLDP advertisements received from other LLDP devices.

Syntax show lldp info remote-device [port-list]

Without the [port-list] option, this command provides a global list of the individual devices it has detected by reading LLDP advertisements. Discovered devices are listed by the inbound port on which they were discovered. Multiple devices listed for a single port indicates that such devices are connected to the switch through a hub.

Discovering the same device on multiple ports indicates that the remote device may be connected to the switch in one of the following ways:

- Through different VLANs using separate links. (This applies to switches that use the same MAC address for all configured VLANs.)*
- Through different links in the same trunk.*
- Through different links using the same VLAN. (In this case, spanning-tree should be invoked to prevent a network topology loop. Note that LLDP packets travel on links that spanning-tree blocks for other traffic types.)*

With the [port-list] option, this command provides a listing of the LLDP data that the switch has detected in advertisements received on the specified ports.

For descriptions of the various types of information displayed by these commands, refer to Table 14-3 on page 14-43.

```
ProCurve# show lldp info remote

LLDP Remote Devices Information

LocalPort | ChassisId | PortId | PortName | SysName
-----+-----+-----+-----+-----
1 | 00 11 85 c6 54 60 | 17 | 17 | HP ProCurve Switch ...
2 | 00 11 85 cf 66 80 | 33 | 33 | HP ProCurve Switch ...
```

Figure 14-23. Example of a Global Listing of Discovered Devices

```
ProCurve(config)# show lldp info remote-device a2

LLDP Remote Device Information Detail

Local Port      : A2
ChassisType     : network-address
ChassisId       : 0f ff 7a 5c
PortType        : mac-address
PortId          : 08 00 0f 14 de f2
SysName         : regDN 3004.<IP-Phone-Data >
System Descr    : regDN 3004.<IP-Phone-Data >,h/w rev 0,ASIC rev 0,f/w Boot FW...
PortDescr       : LAN port

System Capabilities Supported : bridge, telephone
System Capabilities Enabled   : bridge, telephone

Remote Management Address

MED Information Detail
EndpointClass :Class3
Media Policy Vlan id :10
Media Policy Priority :7
Media Policy Dscp    :44
Media Policy Tagged  :False
Poe Device Type      :PD
Power Requested      :47
Power Source         :Unknown
Power Priority        :High
```

Indicates the policy configured on the telephone. A configuration mismatch occurs if the supporting port is configured differently.

Figure 14-24. Example of an LLLDP-MED Listing of an Advertisement Received From an LLDP-MED (VoIP Telephone) Source

Displaying LLDP Statistics

LLDP statistics are available on both a global and a per-port levels. Rebooting the switch resets the LLDP statistics counters to zero. Disabling the transmit and/or receive capability on a port “freezes” the related port counters at their current values.

Syntax show lldp stats [port-list]

The global LLDP statistics command displays an overview of neighbor detection activity on the switch, plus data on the number of frames sent, received, and discarded per-port. The per-port LLDP statistics command enhances the list of per-port statistics provided by the global statistics command with some additional per-port LLDP statistics.

Global LLDP Counters:

Neighbor Entries List Last Updated: *Shows the elapsed time since a neighbor was last added or deleted.*

New Neighbor Entries Count: *Shows the total of new LLDP neighbors detected since the last switch reboot. Disconnecting, then reconnecting a neighbor increments this counter.*

Neighbor Entries Deleted Count: *Shows the number of neighbor deletions from the MIB for AgeOut Count and forced drops for all ports. For example, if the admin status for port on a neighbor device changes from **tx_rx** or **txonly** to **disabled** or **rxonly**, then the neighbor device sends a “shutdown” packet out the port and ceases transmitting LLDP frames out that port. The device receiving the shutdown packet deletes all information about the neighbor received on the applicable inbound port and increments the counter.*

Neighbor Entries Dropped Count: *Shows the number of valid LLDP neighbors the switch detected, but could not add. This can occur, for example, when a new neighbor is detected when the switch is already supporting the maximum number of neighbors. Refer to “Neighbor Maximum” on page 14-78.*

Neighbor Entries AgeOut Count: *Shows the number of LLDP neighbors dropped on all ports due to Time-to-Live expiring.*

— Continued —

— Continued —

Per-Port LLDP Counters:

NumFramesRecvd: Shows the total number of valid, inbound LLDP advertisements received from any neighbor(s) on < port-list >. Where multiple neighbors are connected to a port through a hub, this value is the total number of LLDP advertisements received from all sources.

NumFramesSent: Shows the total number of LLDP advertisements sent from < port-list >.

NumFramesDiscarded: Shows the total number of inbound LLDP advertisements discarded by < port-list >. This can occur, for example, when a new neighbor is detected on the port, but the switch is already supporting the maximum number of neighbors. Refer to “Neighbor Maximum” on page 14-78. This can also be an indication of advertisement formatting problems in the neighbor device.

Frames Invalid: Shows the total number of invalid LLDP advertisements received on the port. An invalid advertisement can be caused by header formatting problems in the neighbor device.

TLVs Unrecognized: Shows the total number of LLDP TLVs received on a port with a type value in the reserved range. This could be caused by a basic management TLV from a later LLDP version than the one currently running on the switch.

TLVs Discarded: Shows the total number of LLDP TLVs discarded for any reason. In this case, the advertisement carrying the TLV may be accepted, but the individual TLV was not usable.

Neighbor Ageouts: Shows the number of LLDP neighbors dropped on the port due to Time-to-Live expiring.

```
ProCurve(config)# show lldp stats

LLDP Device Statistics

Neighbor Entries List Last Updated : 2 hours
New Neighbor Entries Count : 20
Neighbor Entries Deleted Count : 20
Neighbor Entries Dropped Count : 0
Neighbor Entries AgeOut Count : 20

LLDP Port Statistics
```

Port	NumFramesRecvd	NumFramesSent	NumFramesDiscarded
1	628	316	0
2	21	12	0
3	0	252	0
4	446	226	0
5	0	0	0
6	0	0	0
.	.	.	.
.	.	.	.
.	.	.	.

Counters showing frames sent on a port but no frames received on that port indicates an active link with a device that either has LLDP disabled on the link or is not LLDP-aware.

Figure 14-25. Example of a Global LLDP Statistics Display

```
ProCurve(config)# show lldp stats 1

LLDP Port Statistics Detail

PortName : 1
Frames Discarded : 0
Frames Invalid : 0
Frames Received : 658
Frames Sent : 331
TLVs Unrecognized : 0
TLVs Discarded : 0
Neighbor Ageouts : 0
```

Figure 14-26. Example of a Per-Port LLDP Statistics Display

LLDP Operating Notes

Neighbor Maximum. The neighbors table in the switch supports as many neighbors as there are ports on the switch. The switch can support multiple neighbors connected through a hub on a given port, but if the switch neighbor maximum is reached, advertisements from additional neighbors on the same or other ports will not be stored in the neighbors table unless some existing neighbors time-out or are removed.

LLDP Packet Forwarding: An 802.1D-compliant switch does not forward LLDP packets, regardless of whether LLDP is globally enabled or disabled on the switch.

One IP Address Advertisement Per-Port: LLDP advertises only one IP address per-port, even if multiple IP addresses are configured by **lldp config < port-list > ipAddrEnable** on a given port.

802.1Q VLAN Information. LLDP packets do not include 802.1Q header information, and are always handled as untagged packets.

Effect of 802.1X Operation. If 802.1X port security is enabled on a port and a connected device is not authorized, LLDP packets are not transmitted or received on that port. Any neighbor data stored in the neighbor MIB for that port prior to the unauthorized device connection remains in the MIB until it ages out. If an unauthorized device later becomes authorized, LLDP transmit and receive operation resumes.

Neighbor Data Can Remain in the Neighbor Database After the Neighbor Is Disconnected. After disconnecting a neighbor LLDP device from the switch, the neighbor can continue to appear in the switch's neighbor database for an extended period if the neighbor's **holdtime-multiplier** is high; especially if the **refresh-interval** is large. Refer to "Changing the Time-to-Live for Transmitted Advertisements" on page 14-50.

Mandatory TLVs. All mandatory TLVs required for LLDP operation are also mandatory for LLDP-MED operation.

Determining the Switch Port Number Included in Topology Change Notification Traps. Enabling topology change notification on a switch port and then connecting or disconnecting an LLDP-MED endpoint on that port causes the switch to send an SNMP trap to notify the designated management station(s). The port number included in the trap corresponds to the internal number the switch maintains for the designated port, and not the port's external (slot/number) identity. To match the port's external slot/number to the internal port number appearing in an SNMP trap, use the **walkmib ifDescr** command, as shown in the following figure:

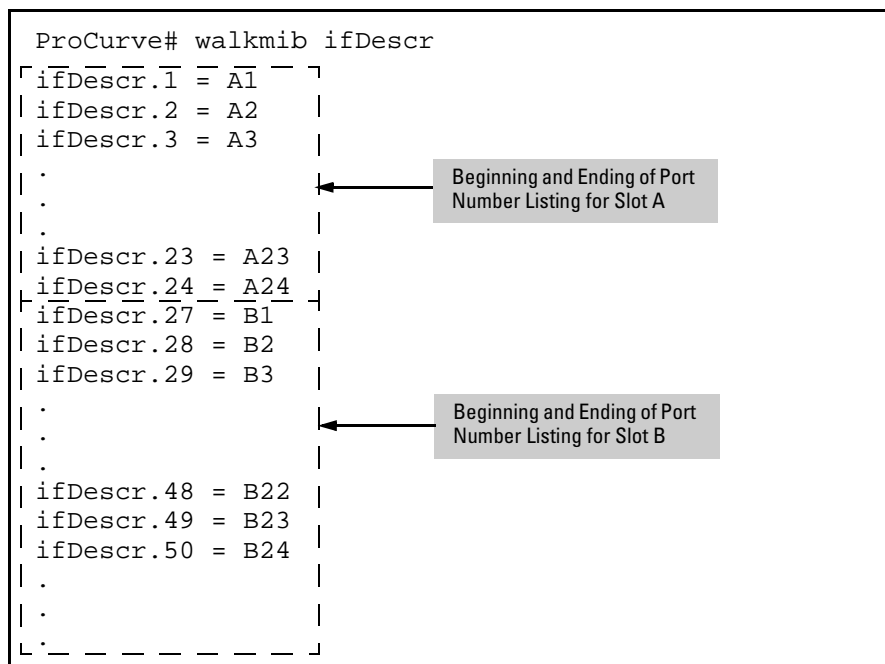


Figure 14-27. Matching Internal Port Numbers to External Slot/Port Numbers

LLDP and CDP Data Management

This section describes points to note regarding LLDP (Link-Layer Discovery Protocol) and CDP (Cisco Discovery Protocol) data received by the switch from other devices. LLDP operation includes both transmitting LLDP packets to neighbor devices and reading LLDP packets received from neighbor devices. CDP operation is limited to reading incoming CDP packets from neighbor devices. (ProCurve switches do not generate CDP packets.)

LLDP and CDP Neighbor Data

With both LLDP and (read-only) CDP enabled on a switch port, the port can read both LLDP and CDP advertisements, and stores the data from both types of advertisements in its neighbor database. (The switch only *stores* CDP data that has a corresponding field in the LLDP neighbor database.) The neighbor database itself can be read by either LLDP or CDP methods or by using the **show lldp** commands. Take note of the following rules and conditions:

- If the switch receives both LLDP and CDP advertisements on the same port from the same neighbor the switch stores this information as two separate entries if the advertisements have differences chassis ID and port ID information.
- If the chassis and port ID information are the same, the switch stores this information as a single entry. That is, LLDP data overwrites the corresponding CDP data in the neighbor database if the chassis and port ID information in the LLDP and CDP advertisements received from the same device is the same.
- Data read from a CDP packet does not support some LLDP fields, such as “System Descr”, “SystemCapSupported”, and “ChassisType”. For such fields, LLDP assigns relevant default values. Also:
 - The LLDP “System Descr” field maps to CDP’s “Version” and “Platform” fields.
 - The switch assigns “ChassisType” and “PortType” fields as “local” for both the LLDP and the CDP advertisements it receives.
 - Both LLDP and CDP support the “System Capability” TLV. However, LLDP differentiates between what a device is capable of supporting and what it is actually supporting, and separates the two types of information into subelements of the System Capability TLV. CDP has only a single field for this data. Thus, when CDP System Capability data is mapped to LLDP, the same value appears in both LLDP System Capability fields.
 - System Name and Port Descr are not communicated by CDP, and thus are not included in the switch’s Neighbors database.

Note

Because ProCurve switches do not generate CDP packets, they are not represented in the CDP data collected by any neighbor devices running CDP.

A switch with CDP disabled forwards the CDP packets it receives from other devices, but does not store the CDP information from these packets in its own MIB.

LLDP data transmission/collection and CDP data collection are both enabled in the switch’s default configuration. In this state, an SNMP network management application designed to discover devices running either CDP or LLDP can retrieve neighbor information from the switch regardless of whether LLDP or CDP is used to collect the device-specific information.

Protocol State	Packet Generation	Inbound Data Management	Inbound Packet Forwarding
CDP Enabled ¹	n/a	Store inbound CDP data.	No forwarding of inbound CDP packets.
CDP Disabled	n/a	No storage of CDP data from neighbor devices.	Floods inbound CDP packets from connected devices to outbound ports.
LLDP Enabled ¹	Generates and transmits LLDP packets out all ports on the switch.	Store inbound LLDP data.	No forwarding of inbound LLDP packets.
LLDP Disabled	No packet generation.	No storage of LLDP data from neighbor devices.	No forwarding of inbound LLDP packets.

¹Both CDP data collection and LLDP transmit/receive are enabled in the default configuration. If a switch receives CDP packets and LLDP packets from the same neighbor device on the same port, it stores and displays the two types of information separately if the chassis and port ID information in the two types of advertisements is different. In this case, if you want to use only one type of data from a neighbor sending both types, disable the unwanted protocol on either the neighbor device or on the switch. However, if the chassis and port ID information in the two types of advertisements is the same, the LLDP information overwrites the CDP data for the same neighbor device on the same port.

CDP Operation and Commands

By default the switches covered in this guide have CDP enabled on each port. This is a read-only capability, meaning that the switch can receive and store information about adjacent CDP devices but does not generate CDP packets.

When a CDP-enabled switch receives a CDP packet from another CDP device, it enters that device's data in the CDP Neighbors table, along with the port number where the data was received (and does not forward the packet). The switch also periodically purges the table of any entries that have expired. (The hold time for any data entry in the switch's CDP Neighbors table is configured in the device transmitting the CDP packet, and cannot be controlled in the switch receiving the packet.) A switch reviews the list of CDP neighbor entries every three seconds, and purges any expired entries.

Command	Page
show cdp	14-83
show cdp neighbors [< port-list > detail] [detail < port-list >]	14-84
[no] cdp run	14-85
[no] cdp enable < port-list >	14-85

Note

For details on how to use an SNMP utility to retrieve information from the switch's CDP Neighbors table maintained in the switch's MIB (Management Information Base), refer to the documentation provided with the particular SNMP utility.

Viewing the Switch's Current CDP Configuration. CDP is shown as enabled/disabled both globally on the switch and on a per-port basis.

Syntax: show cdp

Lists the switch's global and per-port CDP configuration.

The following example shows the default CDP configuration.

```
ProCurve(config)# show cdp
Global CDP information
  Enable CDP [Yes] : Yes

Port CDP
-----
A1  enabled
A2  enabled
A3  enabled
.   .
.   .
.   .
```

Figure 14-28. Example of Show CDP with the Default CDP Configuration

Viewing the Switch's Current CDP Neighbors Table. Devices are listed by the port on which they were detected.

Syntax: show cdp neighbors

Lists the neighboring CDP devices the switch detects, with a subset of the information collected from the device's CDP packet.

[[e] port-numb [detail]]

*Lists the CDP device connected to the specified port. (Allows only one port at a time.) Using **detail** provides a longer list of details on the CDP device the switch detects on the specified port.*

[detail [[e] port-num]]

Provides a list of the details for all of the CDP devices the switch detects. Using port-num produces a list of details for the selected port.

Figure 14-29 lists CDP devices that the switch has detected by receiving their CDP packets.

```
ProCurve> show cdp neighbors
CDP neighbors information
```

Port	Device ID	Platform	Capability
A1	Accounting(0030c1-7fcc40)	J4812A ProCurve Switch...	S
A2	Research(0060b0-889e43)	J4121A ProCurve Switch...	S
A4	Support(0060b0-761a45)	J4121A ProCurve Switch...	S
A7	Marketing(0030c5-38dc59)	J4813A ProCurve Switch...	S
A12	Mgmt NIC(099a05-09df9b)	NIC Model X666	H
A12	Mgmt NIC(099a05-09df11)	NIC Model X666	H

Figure 14-29. Example of CDP Neighbors Table Listing

Enabling CDP Operation. Enabling CDP operation (the default) on the switch causes the switch to add entries to its CDP Neighbors table for any CDP packets it receives from other neighboring CDP devices.

Disabling CDP Operation. Disabling CDP operation clears the switch's CDP Neighbors table and causes the switch to drop inbound CDP packets from other devices without entering the data in the CDP Neighbors table.

Syntax: [no] cdp run

*Enables or disables CDP read-only operation on the switch.
(Default: Enabled)*

For example, to disable CDP read-only on the switch:

```
ProCurve(config)# no cdp run
```

When CDP is disabled:

- **show cdp neighbors** displays an empty CDP Neighbors table
- **show cdp** displays

```
Global CDP information
Enable CDP [Yes]: No
```

Enabling or Disabling CDP Operation on Individual Ports. In the factory-default configuration, the switch has all ports enabled to receive CDP packets. Disabling CDP on a port causes it to drop inbound CDP packets without recording their data in the CDP Neighbors table.

Syntax: [no] cdp enable < [e] port-list >

For example, to disable CDP on port A1:

```
ProCurve(config)# no cdp enable a1
```

Configuring for Network Management Applications
LLDP (Link-Layer Discovery Protocol)

Redundancy (Switches 8200zl)

Contents

Overview	15-2
Terminology	15-2
How the Management Modules Interact	15-3
Using Redundant Management	15-4
Displaying Redundancy Status	15-4
Enabling or Disabling Redundant Management	15-5
Directing the Standby Module to Become Active	15-7
Setting the Active Management Module for Next Boot	15-8
Enabling and Disabling Fabric Modules	15-11
Management Module Switchover	15-12
Events that Cause a Switchover	15-12
When Switchover Will not Occur	15-12
Consequences of Switchover	15-12
Resetting the Management Module	15-13
Hotswapping Management Modules	15-14
Hotswapping Out the Active Management Module	15-14
When the Standby Module is not Available	15-15
Hotswapping In a Management Module	15-15
Software Version Mismatch Between Active and Hotswapped Module	15-15
Downloading a New Software Version	15-16
File Synchronization after Downloading	15-16
Potential Software Version Mismatches After Downloading	15-17
Downloading a Software Version Serially if the Management Module is Corrupted	15-20
Turning Off Redundant Management	15-20

Disabling Redundancy with Two Modules Present	15-20
Disabling Redundancy With Only One Module Present	15-21
Displaying Management Information	15-22
Active Management Module Commands	15-22
Show Modules	15-22
Show Redundancy	15-23
Show Flash	15-24
Show Version	15-24
Show Log	15-25
Standby Management Module Commands	15-26
Show Redundancy	15-26
Show Flash	15-26
Show Version	15-27
Existing CLI Commands Affected by Redundant Management	15-28
Boot Command	15-28
Setting the Default Flash for Boot	15-30
Reload Command	15-31
Additional Commands Affected by Redundant Management	15-33
Using the Web Browser for Redundant Management	15-35
Identity Page	15-35
Overview Page	15-36
Redundancy Status Page	15-36
Device View Page	15-37
Management Module LED Behavior	15-39
Active (Actv) LED Behavior	15-39
Standby Led Behavior	15-39
Logging Messages	15-40
Log File	15-40
Crash Files	15-41
Displaying Saved Crash Information	15-41
Notes on How the Active Module is Determined	15-43
Diagram of Decision Process	15-44
Event Log Messages	15-45

Overview

Redundancy provides the ability to keep your switch operating by using dual management modules, one active module and one standby module. In the event of a failure, the currently active management module will switchover to the standby management module, which then becomes the active management module.

The advantages of redundant management are:

- Maintaining switch operation if a hardware failure occurs on the active management module
- Minimizing restart time because of failure on a management module
- Hotswapping a failed management module with no downtime
- Allowing a faster software upgrade process (less downtime) when updating software versions

Note

The fabric modules are also redundant and can be enabled or disabled. See “Enabling and Disabling Fabric Modules” on page 15-12.

Terminology

Redundant management uses the following terminology.

Active Management Module. A management module that booted successfully and is actively managing the switch.

Standby Management Module. A management module that is ready to become the active management module if the active management module fails.

Failed Management Module. A management module that did not pass selftest and is not in standby mode.

Offline Management Module. A management module that is offline because redundancy is disabled.

Primary Image. The software version stored in primary flash on each management module.

Secondary Image. The software version stored in secondary flash on each management module.

Selftest. A test performed at boot to ensure the management module is functioning correctly. If the module fails selftest, it does not go into active or standby mode. If both modules fail selftest, the switch does not boot.

Switchover. When the other management module becomes the active management module.

How the Management Modules Interact

When the switch boots up, the management modules run selftest to decide which is the active module and which is the standby module (see “Notes on How the Active Module is Determined” on page 15-44). The module that becomes active finishes booting and then brings up the interface modules and ports. The standby module boots to a certain point, syncs basic files such as the config and security files, and only finishes booting if the active management module fails or you choose to change which module is the active module.

The two management modules communicate by sending heartbeats back and forth. The active management module continuously synchronizes the configuration and security files with the standby module. If the active management module fails, the standby management module becomes the active module and finishes the boot process by reading the stored config file, resetting the interface modules, and bringing up the ports.

Note

The management module that becomes the “active” module will be the one that is booted going forward.

Using Redundant Management

There are new CLI commands for redundant management as well as modifications to existing commands. (See “Existing CLI Commands Affected by Redundant Management” on page 15-29)

New Redundant Management Commands	Page
redundancy management-module	below
redundancy switchover	15-8
redundancy active-management	15-9
redundancy fabric-module	15-12
show redundancy	15-5;15-27

Displaying Redundancy Status

You can display the status of both the management and fabric redundant modules using this command:

Syntax: show redundancy

Displays the status of the management and fabric modules.

An example of the output for the show redundancy command is seen in Figure 15-1.

```
ProCurve(config)# show redundancy

Settings
-----
  Mgmt Redundancy : enabled

Statistics
-----
  Failovers      : 0
  Last Failover  :

Slot Module Description                Status  SW Version  Boot Image
-----
1   ProCurve J9092A Management Module 8200zl Active   K.12.XX    Primary
2   ProCurve J9092A Management Module 8200zl Standby  K.12.XX    Primary

1   ProCurve J9093A F2 Fabric Module 8200zl Enabled
2   ProCurve J9093A F2 Fabric Module 8200zl Enabled
```

Figure 15-1. Example of show redundancy Command for Management and Fabric Modules

Enabling or Disabling Redundant Management

You can enable or disable redundant management using this command:

Syntax: [no] redundancy management-module

*Allows enabling or disabling of redundant management. The current active module continues to be the active module on boot unless you use the **redundancy active-management** command to make the other module the active module.*

You are prompted with “All configuration files and software images on the off-line management module will be overwritten with the data from the current active management module. Do you want to continue [y/n]?”

*The **no** version of the command disables redundant management. You are prompted with this message: “The other management module will no longer be used for system redundancy except in the case of a HW failure of the active management module. Do you want to continue [y/n]?”. Selecting “n” disables redundant management.*

The **redundancy management-module** command in Figure 15-2 shows redundant management being enabled. The **show redundancy** command displays “Mgmt Redundancy” as enabled. Management Module 1 is the standby management module and Management Module 2 is the active management module.

```

ProCurve(config)# redundancy management-module
All configuration files and software images on the off-line management
module will be overwritten with the data from the current active
management module. Do you want to continue [y/n]? y
ProCurve(config)#

ProCurve(config)# show redundancy

  Settings
  -----
  Mgmt Redundancy : enabled ← Redundancy enabled

  Statistics
  -----
  Failovers      : 0
  Last Failover  :

Slot Module Description                               Status   SW Version   Boot Image
-----
1   ProCurve J9092A Management Module 8200z1 Standby   K.12.XX     Primary
2   ProCurve J9092A Management Module 8200z1 Active   K.12.XX     Primary

1   ProCurve J9093A F2 Fabric Module 8200z1 Enabled
2   ProCurve J9093A F2 Fabric Module 8200z1 Enabled

```

Figure 15-2. Example of Enabling Redundancy

The **no** version of the **redundancy management-module** command is used to disable redundancy on the switch, as seen in Figure 15-3. The **show redundancy** command displays “Mgmt Redundancy” as disabled. The standby management module in slot 1 is now offline. The management module in slot 2 remains the active management module.

Note

ProCurve recommends that you leave redundancy enabled. If the active management module has a hardware failure, the standby module may take over and may have an old configuration since file synchronization has not occurred.

The **redundancy management-module** command allows you to shut down a management module that is not functioning correctly without physically removing the module. However, removing the management module is the recommended method.

Redundancy (Switches 8200zl) Using Redundant Management

```
ProCurve(config)# no redundancy management-module
The other management module will no longer be used for system
redundancy except in the case of a hardware failure of the active management
module. Do you want to continue[y/n]? y

ProCurve(config)# show redundancy

  Settings
  -----
  Mgmt Redundancy : disabled ← Redundancy disabled

  Statistics
  -----
  Failovers      : 1
  Last Failover  : Tue Mar 19 12:42:31 2007

Slot Module Description                               Status  SW Version  Boot Image
-----
1   ProCurve J9092A Management Module 8200zl  Offline  K.12.XX    Primary
2   ProCurve J9092A Management Module 8200zl  Active   K.12.XX    Primary

1   ProCurve J9093A F2 Fabric Module 8200zl  Enabled
2   ProCurve J9093A F2 Fabric Module 8200zl  Enabled
```

Figure 15-3. Example of Disabling Redundancy

Directing the Standby Module to Become Active

To make the standby management module become the active management module, use the **redundancy switchover** command. The switch will switchover after all files have finished synchronizing. This may take a couple of minutes if there have been recent configuration file changes or if you have downloaded a new operating system. The standby module finishes booting and becomes the active module. The formerly active module becomes the standby module if it passes selftest.

Syntax: redundancy switchover

Causes an immediate switchover to the standby module. The warning displays: “This management module will now reboot and will become the standby module! You will need to use the other management module’s console interface. Do you want to continue [y/n]?”

If redundancy has been disabled, or there is no standby module or the standby module is not in standby mode, this message displays:

The other management module does not exist or is not in standby mode

An example of the **redundancy switchover** command is shown in Figure 15-4.

```
ProCurve(config)# redundancy switchover
This management module will now reboot from primary image and will become
the standby module! You will need to use the other management module's
console interface. Do you want to continue [y/n]? y

ROM information:
  Build directory: /sw/rom/build/bmrom(t2g)
  Build date:      Mar 15 2007
  Build time:      08:24:27
  Build version:   K.12.02
  Build number:    13040
Select profile (primary):

Booting Primary Software Image...
.
.
.

Standby Console#
```

Figure 15-4. An Example of the Redundancy Switchover Command

Setting the Active Management Module for Next Boot

You can select which management module you want to be the active management module at the next “boot system” or switchover event. Enter this command:

Syntax: redundancy active-management <standby | management-module1 | management-module2>

The specified module becomes the active management module at the next system boot. This message displays: “On the next system boot, the <module specified> will become active.”

This command will not take effect if the standby management module has failed selftest.

management-module1: Configures management-module 1 as the active management module for the next system boot.

management-module2: Configures management-module 2 as the active management module for the next system boot.

standby: Configures the current standby module as the active management module for the next system boot if redundancy is enabled. If redundancy is disabled, it will become enabled as a standby module at the next boot or failover event.

Redundancy (Switches 8200z1) Using Redundant Management

If the specified management module is not there or is in failed mode, this message displays:

```
The <specified module> is not present or is in failed state.
```

Figure 15-5 shows an example of setting management module 2 to be the active management module.

```
ProCurve(config)# redundancy active-management management-module2
On the next system boot, the management-module2 will become active.
ProCurve(config)# boot system
/boot occurs...
ProCurve(config)# show redundancy

Settings
-----
Mgmt Redundancy : enabled

Statistics
-----
Failovers       : 0
Last Failover   :

Slot Module Description                               Status  SW Version  Boot Image
-----
1   ProCurve J9092A Management Module 8200z1 Standby  K.12.XX    Primary
2   ProCurve J9092A Management Module 8200z1 Active   K.12.XX    Primary

1   ProCurve J9093A F2 Fabric Module 8200z1 Enabled
2   ProCurve J9093A F2 Fabric Module 8200z1 Enabled
```

Figure 15-5. Setting a Management Module to be Active on the Next Boot

If redundancy has been disabled and you specify the standby module with the **active-management** command, upon rebooting the offline module becomes the standby module. The state of redundancy (enabled or disabled) is based on the value in the configuration file in the offline (now standby) module. The configuration files haven't been synchronized if redundancy has been disabled. An example of making the offline management module become the standby management module when redundancy is disabled is shown in Figure 15-6.


```

ProCurve(config)# show redundancy

Settings
-----
Mgmt Redundancy : Disabled ← Redundancy disabled

Statistics
-----
Failovers      : 0
Last Failover  :

Slot Module Description                Status  SW Version  Boot Image
-----
1   ProCurve J9092A Management Module 8200zl Active   K.12.XX     Primary
2   ProCurve J9092A Management Module 8200zl Offline K.12.XX     Primary

1   ProCurve J9093A Fabric Module 8200zl   Enabled
2   ProCurve J9093A Fabric Module 8200zl   Enabled

ProCurve Switch 8200zl(config)# redundancy active-management standby
On the next system boot, the standby will become active.
Redundancy and Synchronization have been disabled, so it will
not have current configurations.

ProCurve Switch 8200zl(config)# boot
The other management module is not in standby mode and this command will
not cause a switchover. System will reboot from primary image.
Do you want to continue [y/n]? y

(After system reboots...)

ProCurve Switch 8200zl(config)# show redundancy

Settings
-----
Mgmt Redundancy : Disabled ← Redundancy disabled

Statistics
-----
Failovers      : 0
Last Failover  :

Slot Module Description                Status  SW Version  Boot Image
-----
1   ProCurve J9092A Management Module 8200zl Standby  K.12.XX     Primary
2   ProCurve J9092A Management Module 8200zl Active   K.12.XX     Primary

```

Figure 15-6. Example Showing Results of Switching to Standby Module when Redundancy is Disabled

Enabling and Disabling Fabric Modules

The fabric modules can be enabled or disabled even if they are not present in the switch. You cannot disable both fabric modules at the same time; one must be enabled. Use this command to enable or disable the redundant fabric modules. Disabling one fabric module reduces the overall switching capacity of the 8200zl series switches. On some networks where network utilization is less than 50%, you may not notice any degradation of performance.

Syntax: redundancy fabric-module [1 | 2] [enable | disable]

Allows enabling or disabling of fabric modules. You cannot have both fabric modules disabled at the same time.

Default: Both fabric modules are enabled.

```
ProCurve(config)# redundancy fabric-module 2 disable
ProCurve(config)# show redundancy

Settings
-----
Mgmt Redundancy : enabled

Statistics
-----
Failovers       : 0
Last Failover   :

Slot Module Description                Status  SW Version  Boot Image
-----
1   ProCurve J9092A Management Module 8200zl Active   K.12.XX     Primary
2   ProCurve J9092A Management Module 8200zl Standby  K.12.XX     Primary

1   ProCurve J9093A F2 Fabric Module 8200zl Enabled
2   ProCurve J9093A F2 Fabric Module 8200zl Disabled
```

Figure 15-7. Example of Disabling a Fabric Module

Management Module Switchover

Events that Cause a Switchover

There are a number of events that can cause the active management module to switchover to the standby management module when redundancy is enabled:

- The active management module crashes
- The standby management module does not receive a heartbeat from the active management module
- The **redundancy switchover** command is executed
- The active management module is hotswapped out
- The **MM Reset** button on the active management module is pressed
- The **MM Shutdown** button on the active management module is pressed
- The **boot** or **boot active** command is executed
- The **reload** command is executed
- There is a hardware failure on the active management module

In all of these cases the standby management module takes control and performs the actual switchover. The reason for the switchover is entered in log messages on the newly active management module and to any configured Syslog servers.

When Switchover Will not Occur

There are some events for which a switchover is not triggered:

- When a **boot system** command is executed
- When the **Clear** button on the System Support module is pressed
- When redundancy is disabled, unless there is a hardware failure and the system is rebooted.

Consequences of Switchover

When a switchover occurs, the standby management module completes its boot process and reloads all the interface modules. The following information is not saved:

- Port Counter information (counters are reset)

- Learned routes (from routing protocols)
- MAC addresses
- IGMP, LACP, GVRP, LLDP, CDP, 802.1X, STP, VRRP, PIM learned data
- Web auth and MAC auth connections
- IDM data
- AAA accounting data
- Telnet connection to the switch
- SNMP sample rates

Resetting the Management Module

The **MM Reset** button found on each management module reboots its management module. If the management module is active and redundancy is enabled, switchover occurs. The standby management module is notified immediately. It then takes over and becomes the active management module. If the **MM Reset** button is pressed on the standby management module, that module reboots but no other switch operations are affected. The active management module remains in control.

If redundancy is disabled, the active management module reboots and remains in control, as long as it passes selftest.

Caution

ProCurve does not recommend using the MM Reset button to trigger a switchover. Files being copied over at the time of the reset will be aborted.

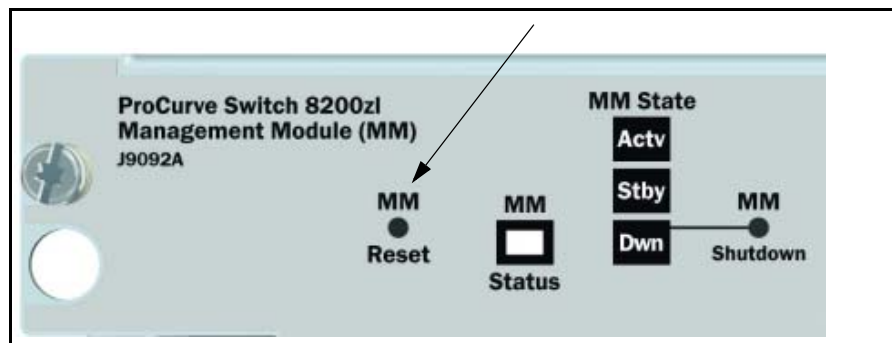


Figure 15-8. The MM Reset Button on the 8200zl Management Module

Hotswapping Management Modules

Hotswapping Out the Active Management Module

You can hotswap out the active management module and have switch operations taken over by the standby management module by following the correct shutdown procedure on the active module using the **MM Shutdown** button. When the **MM Shutdown** button is pressed, any file synchronization in progress completes before the shutdown begins, and then a graceful shutdown of that management module occurs.

1. On the management module to be hotswapped out, press the **MM Shutdown** button. It is located between the Module Operation and Component Status LEDs. See Figure 15-9.




Figure 15-9. The MM Shutdown Button

2. The **Dwn** LED to the right of the MM Shutdown button will begin flashing green. File synchronization will complete before shutdown occurs.
3. The standby module takes control and the switchover occurs. It is now the active management module.
4. The **Dwn** LED on the management module being hotswapped out turns green and all other LEDs go out when it is OK to remove the module.
5. The module being hotswapped out goes into offline mode. In the “offline” mode, the module cannot take over when the active module fails over.

When the Standby Module is not Available

If you have disabled redundancy with the **no redundancy management-module** command, or the standby module failed selftest, the **Down** LED will not turn green to indicate it is OK to hotswap out the active management module.

Note

If you remove the active management module without pressing the **MM Shutdown** button, any files that may have been in the process of synchronizing will not finish synchronizing to the standby module and all file transfer is aborted.

Hotswapping In a Management Module

If another management module is hotswapped in while there is an active management module booted up, the newly hotswapped management module becomes the standby module. It partially boots up and heartbeats are sent back and forth with the active management module. No negotiating is needed as to which module will become the active management module as there is already a functioning active management module. However, these conditions must be met to determine if the hotswapped module can become a standby management module:

- The hotswapped module must pass selftest
- Redundancy is not administratively disabled (using the **no redundancy management-module** command). If the active management module's config file has redundancy administratively disabled, the hotswapped management module goes into "offline" mode.

Software Version Mismatch Between Active and Hotswapped Module

Sometimes the software version in the hotswapped module may not match the software version in the active module. In these cases the following occurs:

1. The active module sends the primary and secondary images in flash to the hotswapped module.
2. The module that was hotswapped in then reboots if necessary to primary or secondary flash, whichever matches (if it doesn't match already).
3. After the hotswapped management module finishes booting, it is sent the config and other critical files from the active management module.
4. The hotswapped management module goes into standby mode and is ready to take over in case of a switchover.

Downloading a New Software Version

File Synchronization after Downloading

After downloading a new software version to either the primary or secondary flash of the active management module, the software version is immediately copied to the corresponding flash (primary or secondary) of the standby module unless the standby module failed selftest or redundancy was disabled with the **no redundancy management-module** command.

The configuration files, including which configuration file to use for that flash image, are synchronized. For example, if the active management module is using config1, the standby module will also be synchronized to use config1.

Table 15-1. Example of Upgrading Software Version K.12.03 to Version K.12.04

	Newer Code to Secondary Flash		New Code to Primary Flash	
	Active MM	Standby MM	Active MM	Standby MM
Software version downloaded to Primary flash image	K.12.03	K.12.03	K.12.04	K.12.04
Software version downloaded to Secondary flash image	K.12.04	K.12.04	K.12.03	K.12.03

Note

See “Setting the Default Flash for Boot” on page 15-31 for information about testing new software versions.

After installing the new software to the active management module, wait a few minutes, and then verify that the standby module has been synchronized with the new software as well (use the **show flash** command). If the default flash for boot has been set correctly, you can start the standby management module on the new software by executing the **boot standby** command. This does not interrupt current switch operations yet. After the standby management module has rebooted and is ready for takeover in standby mode

(you can verify this using the **show redundancy** command), you can now switch over to the management module running the newer software with this command:

```
ProCurve# redundancy switchover
```

This causes a switchover to the management module that received the new software version, which becomes the active management module. This method incurs the least amount of network downtime for booting. If downtime is not an issue, use the **boot system** command. Both management modules will then be running the new software version.

Potential Software Version Mismatches After Downloading

When a new software version is downloaded to the active management module, it is immediately copied to the corresponding flash (primary or secondary) in the standby management module unless redundancy has been disabled. If the standby management module is rebooted, it will be running a different software version than the active management module. You can direct the standby module to boot from the non-corresponding flash image that has a different software version during the actual reboot process of the standby module when the prompt to select the Boot Profile appears.

```
Standby Console# show flash
Image           Size(Bytes)   Date   Version   Build #
-----
Primary Image   : 7493854   03/21/07 K.12.XX   1617
Secondary Image : 7463821   03/05/07 K.12.XX   351
Boot Rom Version: K.12.03
Default Boot    : Primary

Boot Profiles:
0. Monitor ROM Console
1. Primary Software Image
2. Secondary Software Image

Select profile(primary): 2
Booting Secondary Software Image
```

The diagram shows a terminal window with a table of flash images and a list of boot profiles. Two callout boxes provide context: one points to the 'Boot Profiles' list, stating 'You can select which flash to boot from at this point in the boot process.', and another points to the '2' in the 'Select profile(primary): 2' line, stating 'Indicates the default boot choice'. The terminal output shows that the secondary software image is being booted.

Figure 15-10. Booting the Standby Management Module to Secondary Flash

Caution

If you have booted one module out of primary flash and one module out of secondary flash, and the secondary flash is running a prior software version because the latest version was never copied over from the primary flash, you will have a software version mismatch. The configuration file may not work with that software version. See “Software Version Mismatch Between Active and Hotswapped Module” on page 15-16 for more information.

Additionally, if a switchover occurs, or if you reboot to make the standby module become the active module, any configuration file changes made may not work on the active module if it has a different software version from the standby module.

When you enter the **show redundancy** command and a software version mismatch exists, a warning message is displayed, as shown at the bottom of Figure 15-11.

Redundancy (Switches 8200z1)
Downloading a New Software Version

```
ProCurve(config)# show version
Management Module 1: Active
Image stamp:      /sw/code/build/btm(t2g)
                  Mar 15 2007 12:28:32
                  K.12.30
                  64
Boot Image:      Primary

Management Module 2: Standby
Image stamp:      /sw/code/build/btm(t2g)
                  Mar 21 2007 14:24:38
                  K.12.30
                  789
Boot Image:      Secondary

ProCurve(config)# show redundancy

Settings
-----
Mgmt Redundancy : Enabled

Statistics
-----
Failovers       : 0
Last Failover   :

Slot Module Description                               Status  SW Version  Boot Image
-----
1   ProCurve J9092A Management Module 8200z1  Active  K.12.30     Primary
2   ProCurve J9092A Management Module 8200z1  Standby K.12.30     Primary

1   ProCurve J9093A F2 Fabric Module 8200z1  Enabled
2   ProCurve J9093A F2 Fabric Module 8200z1  Enabled

Warning: Standby module is running a different software version and may be using
a different configuration file. Configuration changes on active management
module may not take effect on a failover.
```

Mismatch exists

Figure 15-11. Example of a Software Version Mismatch Between the Active and Standby Modules

Downloading a Software Version Serially if the Management Module is Corrupted

If the software version on a management module becomes corrupted, you may need to do a serial download to restore the affected module. The non-corrupted management module becomes the active module. You can then use the serial port on the corrupted management module to download a new software version. When the corrupted module is rebooted, the software version in the corrupted module is immediately overwritten by the software version in the active management module. Both management modules should now be operating on the same software version.

Turning Off Redundant Management

Disabling Redundancy with Two Modules Present

In some cases, for troubleshooting a suspect management module you may want to operate the switch with redundant management disabled by entering this command:

```
ProCurve(config)# no redundancy management-module
```

After executing this command, the second management module will not boot into standby mode; it is off line and no longer receives configuration file changes from the active module. The active management module updates its config file with the information that redundancy is disabled.

Note

Even if redundancy has been disabled, the specified management module will become the active management module at the next system boot if you use the **redundancy active-management** command. You are warned that you may not be using current configurations. See “Setting the Active Management Module for Next Boot” on page 15-9.

The second management module is enabled as the active management module in the event of a hardware failure of the first management module.

Figure 15-12 shows that redundant management was disabled.

```
ProCurve(config)# no redundancy management-module
The other management module will no longer be used for system
redundancy except in the case of a hardware failure of the active
management module. Are you sure [y/n]? y

ProCurve(config)# show redundancy

Settings
-----
  Mgmt Redundancy : disabled

Statistics
-----
  Failovers       : 0
  Last Failover   :

Slot Module Description                               Status  SW Version  Boot Image
-----
1   ProCurve J9092A Management Module 8200z1  Offline  K.12.XX    Primary
2   ProCurve J9092A Management Module 8200z1  Active   K.12.XX    Primary

1   ProCurve J9093A F2 Fabric Module 8200z1  Enabled
2   ProCurve J9093A F2 Fabric Module 8200z1  Enabled
```

Figure 15-12. Results of Disabling Redundancy

Disabling Redundancy With Only One Module Present

If you disable redundancy when there is only one management module in the switch, and then you insert a second management module, the second module will never go into standby mode. You must re-enable redundant management using this command:

```
ProCurve(config)# redundancy management-module
```

The currently active module remains active on boot (assuming no selftest failure) unless you make the newly inserted management module active using this command:

```
ProCurve(config)# redundancy active-management  
standby
```

The standby management module becomes the active management module.

Displaying Management Information

Active Management Module Commands

Show Modules

The **show modules** command displays information about all the modules in the switch as well as additional component information for the following:

- System Support Modules (SSM)—identification, including serial number
- Mini-GBICS—a list of installed mini-GBICs displaying the type, “J” number, and serial number (when available)

Syntax: show modules [details]

Displays information about the installed modules, including:

- *The slot in which the module is installed*
- *The module description*
- *The serial number*
- *The System Support Module description, serial number, and status (8200zl switches only)*

Additionally, the part number (J number) and serial number of the chassis is displayed.

Redundancy (Switches 8200zl)
 Displaying Management Information

```
ProCurve(config)# show modules details

Status and Counters - Module Information

Chassis: 8212zl J8715A          Serial Number:  SG560TN124
Slot  Module Description          Serial Number  Status
-----
MM1   ProCurve J9092A Management Module 8200zl  AD722BX88F    Active
SSM   ProCurve J8784A System Support Module  AF988DC78G    Active
C     ProCurve J8750A 20p +4 Mini-GBIC Module  446S2BX007    Active
      GBIC 1: J4859B 1GB LX-LC             4720347DFED734
      GBIC 2: J4859B 1GB LX-LC             4720347DFED735
```

Figure 15-13. An Example of the show modules details Command for the 8212zl Showing SSM and Mini-GBIC Information

Show Redundancy

The **show redundancy** command displays information about the management and fabric modules. It displays the flash image last booted from, even if the **boot set-default** command has been set to change the flash booted from on the next boot.

```
ProCurve(config)# show redundancy

Settings
-----
Mgmt Redundancy : enabled

Statistics
-----
Failovers      : 0
Last Failover  :

Slot Module Description          Status  SW Version  Boot Image
-----
1    ProCurve J9092A Management Module 8200zl  Standby  K.12.XX    Primary
2    ProCurve J9092A Management Module 8200zl  Active   K.12.XX    Secondary

1    ProCurve J9093A F2 Fabric Module 8200zl  Enabled
2    ProCurve J9093A F2 Fabric Module 8200zl  Enabled
```

The active management module was last booted from secondary flash. The standby management module was last booted from primary flash.



Figure 15-14. Example of show redundancy Command

Show Flash

The **show flash** command displays which software version is in each flash image. The Default Boot field displays which flash image will be used for the next boot.

```
ProCurve(config)# show flash
Image           Size(Bytes)   Date    Version   Build #
-----
Primary Image   : 7463821   03/05/07 K.12.XX   351
Secondary Image : 7463821   03/05/07 K.12.XX   351
Boot Rom Version: K.12.01
Default Boot    : Primary
```

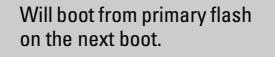


Figure 15-15. Example of Show Flash Command

Show Version

The **show version** command displays system software image information for both management modules as well as which module is the active management module and which is the standby management module. The Boot Image field displays which flash image last booted from, even if the **boot set-default** command has been set to change the flash booted from on the next boot. The output of the **show version** command when redundancy is enabled is shown in Figure 15-16.

```
ProCurve(config)# show version
Management Module 1: Standby
Image stamp:      /sw/code/build/btm(t2g)
                  Mar 5 2007 13:20:59
                  K.12.XX
                  351
Boot Image:       Primary
Management Module 2: Active
Image stamp:      /sw/code/build/btm(t2g)
                  Mar 5 2007 13:20:59
                  K.12.XX
                  351
Boot Image:       Primary
```

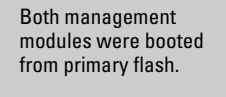


Figure 15-16. Example of Show Version Command when Redundancy is Enabled

When redundancy is disabled, the output of the **show version** command changes, as shown in Figure 15-17.

Redundancy (Switches 8200zl)

Displaying Management Information

```
ProCurve(config)# show version
Management Module 1: Redundancy and Synchronization has been disabled:
                    enable with the 'redundancy' command.

Management Module 2: Active
Image stamp:       /sw/code/build/btm(t2g)
                  Mar  5 2007 13:20:59
                  K.12.XX
                  351
Boot Image:       Primary
```

Figure 15-17. Example of show version Command when Redundancy is Disabled

Show Log

The **show log** command displays the status of the switch and its management modules. See “Logging Messages” on page 15-41. To show log messages in reverse chronological order (most recent messages displayed first), enter **show log -r**.

```
ProCurve Switch 8200zl(config)# show log
Keys:   W=Warning   I=Information
        M=Major     D=Debug
----  Event Log listing: Events Since Boot  ----
M 01/26/14 17:34:07 sys: 'System reboot due to Power Failure'
I 01/26/14 17:34:07 00061 system: -----
--
I 01/26/14 17:34:07 00062 system: Mgmt Module 2 went down without saving crash
information
I 01/26/14 17:36:14 00264 system: Mgmt Module 1 Failed Selftest
I 01/26/14 17:36:19 00068 chassis: Fabric 1 Inserted
I 01/26/14 17:36:19 00068 chassis: Fabric 2 Inserted
I 01/26/14 17:36:19 00068 chassis: Slot D Inserted
I 01/26/14 17:36:19 00690 udpf: DHCP relay agent feature enabled
I 01/26/14 17:36:19 00400 stack: Stack Protocol disabled
I 01/26/14 17:36:19 00128 tftp: Enable succeeded
I 01/26/14 17:36:19 00417 cdp: CDP enabled
I 01/26/14 17:36:19 00688 lldp: LLDP - enabled
I 01/26/14 17:36:19 00066 system: Mgmt Module 2 Booted
I 01/26/14 17:36:19 00260 system: Mgmt Module 2 Active
I 01/26/14 17:36:19 00066 system: Mgmt Module 1 Booted
I 01/26/14 17:36:19 00261 system: Mgmt Module 1 in Standby Mode
```

Figure 15-18. An Example of the Show Log Command Output

Standby Management Module Commands

The standby management module, by design, has very little console capability. You can use three commands—**show flash**, **show version**, and **show redundancy**. The **show redundancy** command displays when a management module is in standby mode.

Show Redundancy

Use the **show redundancy** command to display redundancy status on the standby module, as shown in Figure 15-19. It displays the flash image last booted from, even if the **boot set-default** command has been set to change the flash booted from on the next boot.

```
Standby Console> show redundancy

Settings
-----
  Mgmt Redundancy : Enabled

Statistics
-----
  Failovers       : 1
  Last Failover  : Mon Sep 26 09:50:40 2005

Slot Module Description                               Status  SW Version  Boot Image
-----
1   ProCurve J9092A Management Module 8200z1  Active   K.12.XX    Secondary
2   ProCurve J9092A Management Module 8200z1  Standby  K.12.XX    Primary

1   ProCurve J9093A F1 Fabric Module 8200z1  Enabled
2   ProCurve J9093A F2 Fabric Module 8200z1  Enabled
```

The active management module was last booted from secondary flash. The standby management module was last booted from primary flash.

Figure 15-19. Example of Show Redundancy Command for Standby Module

Show Flash

You can display the flash information on the standby module, as shown in Figure 15-20. The Default Boot field displays which flash image will be used for the next boot.

```
Standby Console> show flash
Image                Size(Bytes)   Date   Version  Build #
-----
Primary Image       : 7493854   03/21/07 K.12.XX  1617
Secondary Image     : 7463821   03/05/07 K.12.XX   351

Boot Rom Version: K.12.03
Default Boot     : Primary
```

Will boot from primary flash on the next boot.

Figure 15-20. Example of Show Flash Command for Standby Module

Show Version

You can display the version information on the standby module, as shown in Figure 15-21. The Boot Image field displays which flash image was last booted from, even if the **boot set-default** command has been set to change the flash booted from on the next boot. Unlike executing the **show version** command on an active management module, this only shows the running version of software on the standby management module.

```
Standby Console> show version
Image stamp:      /sw/code/build/btm(t2g)
                  Mar 21 2007 15:03:31
                  K.12.XX
                  1617
Boot Image:       Primary
```

Was booted from primary flash.

Figure 15-21. Example of Show Version Command for Standby Module

Existing CLI Commands Affected by Redundant Management

Several existing commands have changes related to redundant management.

Boot Command

In redundant management systems, the **boot** or **boot active** command causes a switchover to the standby management module as long as the standby module is in standby mode. This message displays:

```
This management module will now reboot and will become
the standby module! You will need to use the other
management module's console interface. Do you want to
continue [y/n]?
```

If you select “y”, switchover is initiated by the standby management module, which becomes the active management module after boot completes.

If the standby module is not in standby mode (for example, it is in failed mode or offline mode), switchover to the standby module does not occur. The system is rebooted. This message displays:

```
The other management module is not in standby mode and
this command will not cause a switchover, but will
reboot the system, do you want to continue [y/n]?
```

If the other management module is not present in the switch, the system simply reboots.

The **boot** command has these options.

Command	Action
Boot <cr>	Reboots the active management module from the flash image that is specified for the default boot. This can be changed with the boot set-default flash command. You can select which image to boot from during the boot process itself. See Figure 15-22. The switch will switchover to the standby management module. Note: This is changed from always booting from primary flash. You are prompted with a message which will indicate the flash being booted from.

Redundancy (Switches 8200zl)

Existing CLI Commands Affected by Redundant Management

Command	Action
Boot active	Boots the active management module. The switch starts to boot from the default flash image. You can select which image to boot from during the boot process itself. See Figure 15-22. The switch will switchover to the standby management module. If a second management module is not present in the switch, the system is rebooted.
Boot standby	Boots the standby management module. The switch does not switchover. If the standby module is not present, this message displays: "The other management module is not present."
boot system [flash <primary secondary>]	Boots both the active and standby management modules. You can specify the flash image to boot from.
boot set-default flash <primary secondary>	Sets the default flash for the next boot to primary or secondary. You will see this message: "This command changes the location of the default boot. This command will change the default flash image to boot from <flash chosen>. Hereafter, 'reload' and 'boot' commands will boot from <flash chosen>. Do you want to continue [y/n]?"

You can select a boot profile during the reboot process, as shown in Figure 15-22. If you make no selection, the boot defaults to the imaged displayed as the default choice (shown in parentheses).

```
Boot Profiles:
0. Monitor ROM Console
1. Primary Software Image
2. Secondary Software Image

Select profile(primary): 2

Booting Secondary Software Image...
```

Figure 15-22. The Management Module Rebooting, Showing Boot Profiles to Select

An example of the **boot** command with the default flash set to secondary is shown in Figure 15-23.

```
ProCurve(config)# boot set-default flash secondary
This command changes the location of the default boot. This command will
change the default flash image to boot from secondary. Hereafter,
'reload' and 'boot' commands will boot from secondary. Do you want to
continue [y/n]? y

ProCurve(config)# show flash
Image                Size(Bytes)   Date    Version  Build #
-----
Primary Image       : 7476770    03/15/07 K.12.XX    64
Secondary Image    : 7476770    03/15/07 K.12.XX    64
Boot Rom Version: K.12.02
Default Boot       : Secondary

ProCurve(config)# boot
This management module will now reboot from secondary and will become
the standby module! You will need to use the other management module's
console interface. Do you want to continue [y/n]?
```

Figure 15-23. Example Showing boot Command with Default Flash set to Secondary

Caution

For a given reboot, the switch automatically reboots from the startup-config file assigned to the flash (primary or secondary) being used for the current reboot. The **startup-default** command can be used to set a boot configuration policy. This means that both the flash image and one of the three configuration files can be specified as the default boot policy. For more information on multiple configuration files and how they are used, see “*Multiple Configuration Files*” in the “*Switch Memory and Configuration*” chapter in this guide.

Setting the Default Flash for Boot

You can set which flash image to boot from as the default image on boot by using this command:

Syntax: boot set-default flash <primary | secondary>

Sets the flash image to boot from on the next boot.

primary: Boots the primary flash image.

secondary: Boots the secondary flash image.

Figure 15-24 shows an example of the output when the command is used to set the boot default to secondary flash.

```
ProCurve(config)# show flash
Image           Size(Bytes)   Date   Version   Build #
-----
Primary Image   : 7463821   03/05/07 K.12.XX   351
Secondary Image : 7463821   03/05/07 K.12.XX   351
Boot Rom Version: K.12.01
Default Boot    : Primary

ProCurve(config)# boot set-default flash secondary
This command changes the location of the default boot. This
command will change the default flash image to boot from
secondary. Hereafter, 'reload' and 'boot' commands will boot
from secondary. Do you want to continue [y/n]? y

ProCurve(config)# show flash
Image           Size(Bytes)   Date   Version   Build #
-----
Primary Image   : 7463821   03/05/07 K.12.XX   351
Secondary Image : 7463821   03/05/07 K.12.XX   351
Boot Rom Version: K.12.01
Default Boot    : Secondary
```

Figure 15-24. Example of boot set-default Command Defaulting to Secondary Flash

Reload Command

The **reload** command boots the active management module from the current default flash (You can change the default flash with the **boot set-default** command. See “Setting the Default Flash for Boot” on page 15-31). Switchover occurs if redundancy is enabled and the standby management module is in standby mode. If redundancy is disabled or the standby management module is not present, the **reload** command boots the system.

Note

The reload command is a “warm” reboot; it skips the Power on Self Test routine.

Command	Action
reload <cr>	<p>Boots (warm reboot) the active management module. Switchover to the standby management module occurs if redundancy is enabled. If redundancy is disabled or there is no standby management module, the reload command boots the system.</p> <p>Note: If the running config file is different from the stored config file, you will be prompted to save the config file. The reload at/after versions of this command do not display a prompt to save configuration file changes; the changes are lost on the scheduled reload.</p>

```

ProCurve(config)# reload
This command will cause a switchover to the other management module
which may not be running the same software image and configurations.
Do you want to continue [y/n]? y

(Boots...)

ProCurve(config)# show redundancy

  Settings
  -----
  Mgmt Redundancy : Enabled

  Statistics
  -----
  Failovers       : 1
  Last Failover  : Mon April 30 09:10:11 2007

Slot Module Description                               Status   SW Version  Boot Image
-----
1   ProCurve J9092A Management Module 8200zl  Active    K.12.XX    Primary
2   ProCurve J9092A Management Module 8200zl  Standby   K.12.XX    Primary

```

Figure 15-25. Example of Reload Command with Redundancy Enabled

Additional Commands Affected by Redundant Management

The other existing commands operate with redundant management as shown below.

Command	Action
auto-tftp	If a new image is downloaded using auto-tftp , the active management module downloads the new software version to both the active and standby modules. Rebooting after the auto-tftp completes reboots the entire system.
banner	The banner will not be seen on the standby module, only the active module.
chassislocate	If the management module performs a switchover, the LED does not remain lit.
clear	The clear crypto command causes public keys to be deleted from both modules when the second module is in standby mode.
console	Console settings, such as mode, flow-control, and baud-rate, are the same on both management modules. There cannot be individual settings for each management module.
copy	Files are automatically sync'd from the active management module to the standby management module. When no parameter is specified with the copy crash-data or copy crash-log command, files from all modules (management and interface) are concatenated. See "Crash Files" on page 15-42. Note: If redundancy is disabled or the standby module failed selftest, the copy command affects only the active management module.
crypto	Authentication files for ssh or the https server are copied to the standby management module. The clear crypto command deletes the public keys from both modules when the second module is in standby mode.
erase flash	Erases the software version on the active and standby modules. If redundancy has been disabled, or the standby module has not passed selftest, the flash is not erased on the standby module.
erase config	Erases the config file on the active and standby modules. If redundancy has been disabled, or the standby module has not passed selftest, the config file is not erased on the standby module.
erase startup-config	Affects both modules if the second module is in standby mode. If redundancy has been disabled, or the standby module has not passed selftest, the startup-config file is not erased on the standby module.

Command	Action
fastboot	When fastboot is enabled, this information is saved to the standby management module when the config files are sync'd. The fastboot value is used during the next boot on both modules.
front-panel-security factory-reset password-clear password-recovery	This command and its options only affects the active management module. See the section on "Front-Panel Button Functions" in the <i>Access Security Guide</i> for more information about resetting the switch.
kill	Does not affect the console on the standby module.
log	Log messages from a formerly active management module are available on the current active management module after a switchover.
password (set or clear)	Affects only the active management module until a switchover occurs, at which time it affects the new active module.
startup-default	Affects both modules. The config file is immediately sent to the standby module and also becomes the default on that module when the next boot occurs.
update	Only affects the active module. The standby may become the active module when the updated active module is booted.
write	A write memory updates the config file in flash on the active module. The file is then sync'd to the standby module.

Using the Web Browser for Redundant Management

The web browser interface can be used to display information about the active and standby management modules. To learn more about using the web browser interface on your switch, see the chapter “Using the ProCurve Web Browser Interface” in this guide.

Online Help is available for the web browser interface. You can use it by clicking on the question mark button in the upper right corner of any of the web browser interface screens.

Identity Page

The Identity page displays information about the version of software running on both the active and the standby management module.

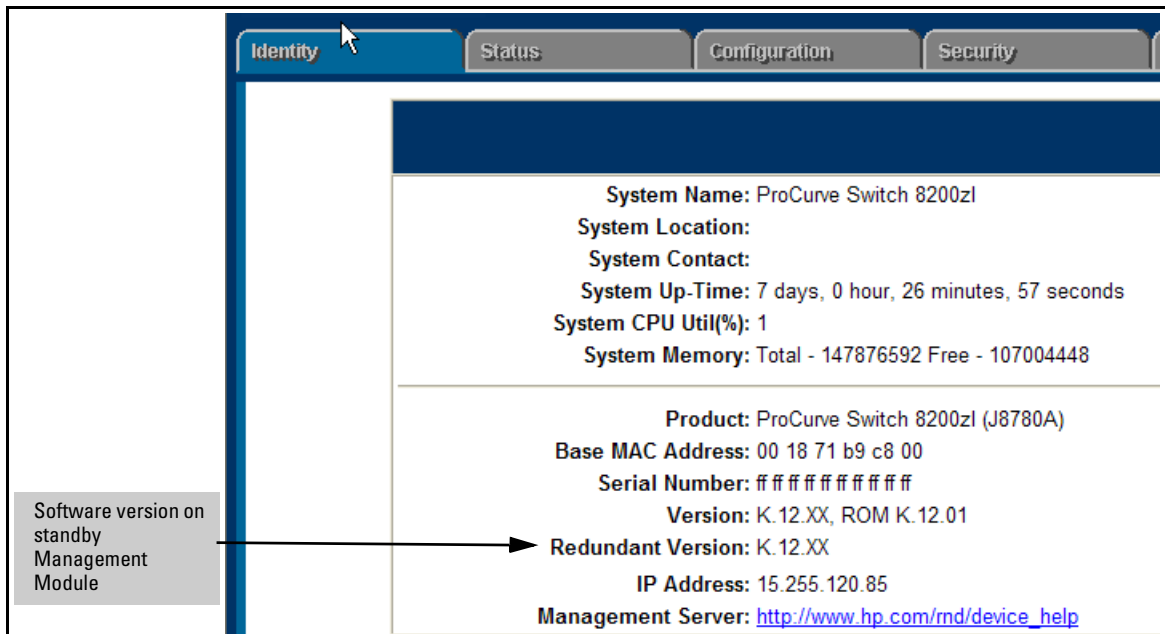


Figure 15-26. Identity Page showing Software Version on both Active and Redundant Management Modules

Overview Page

To view status information about the management modules select the **Status** tab, and then the **Overview** button. The following information is shown:

- Which module is the active module and which is the standby module
- Version of software running on each management module
- The SystemUp Time since the last reboot.

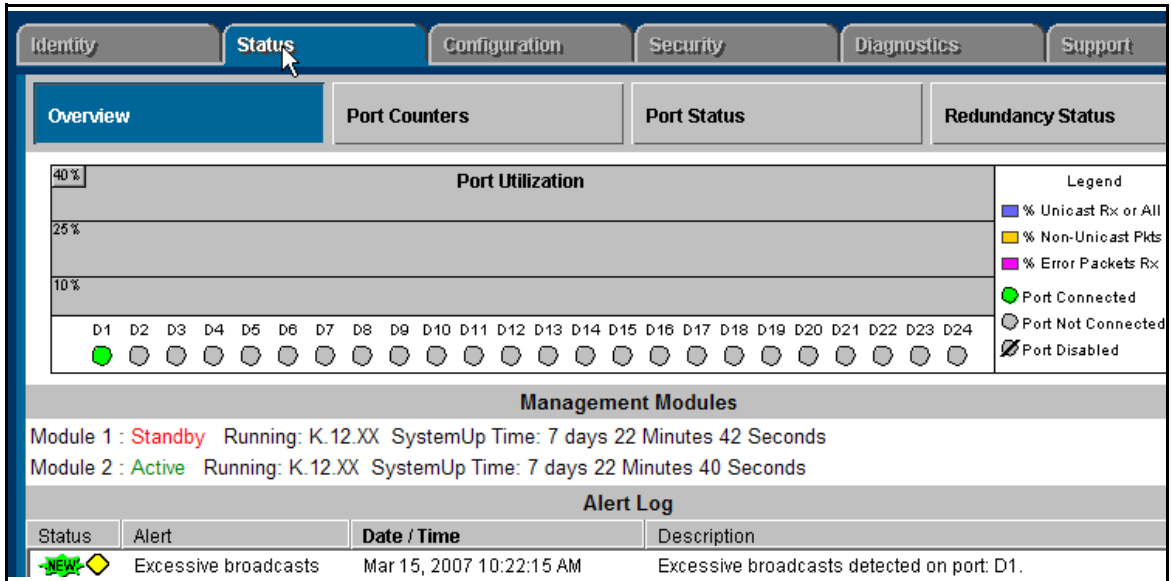


Figure 15-27. Overview Page Showing the SystemUp Time for Both Management Modules

Redundancy Status Page

The **Redundancy Status** tab is visible only if the alternate management module (non-active module) is in standby mode. Select the **Status** tab and then the **Redundancy Status** button. The **Redundancy Status** page displays information about the active and standby management modules and the two fabric modules.

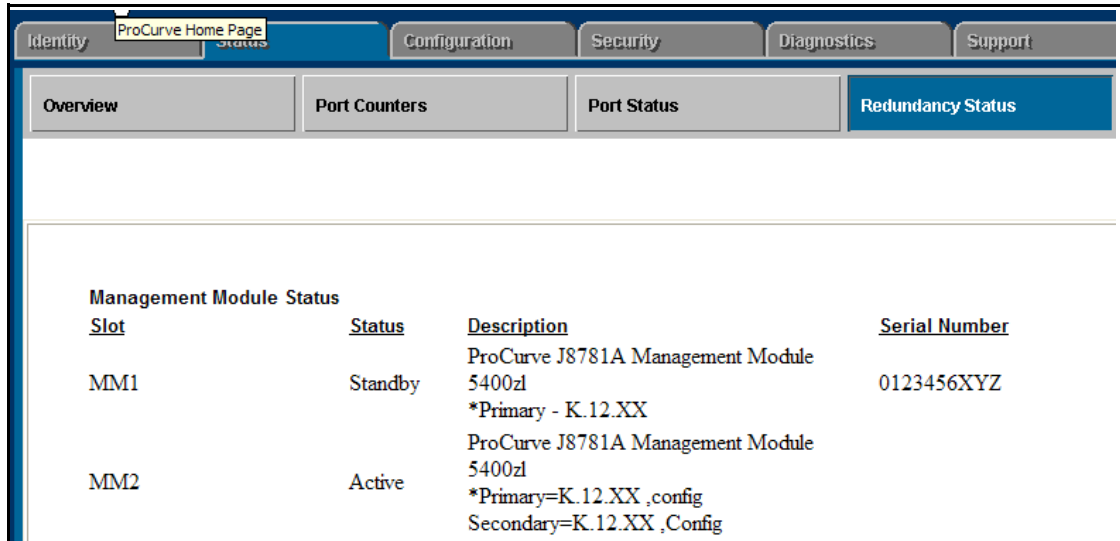


Figure 15-28.Redundancy Status Page Showing Information about the Active and Standby Modules

Device View Page

The **Device View** page displays a graphical representation of the switch. Select the **Configuration** tab and then the **Device View** button. The information displayed includes:

- Fabric modules
- Interface modules
- System Support module
- LEDs and the status of the switch and management modules

The LEDs indicate in green which management module is active and which management module is in standby mode.




Figure 15-29. Device View Showing Two Management Modules

Management Module LED Behavior

Active (Actv) LED Behavior

The Actv (Active) LED shows the LED behavior for various states on the active and standby management modules. See Table 15-2 for the available states and what they indicate. Refer to the *Installation and Getting Started Guide* for your switch for more information about LEDs.




Figure 15-30. The Actv LED on the Management Module

Table 15-2. Actv (Active) LED Behavior for Management Modules

Active Module	Standby Module	Meaning
Solid green	Not lit	Correct Active/Standby mode
Solid green	fast orange flashing	Failed standby module
Not lit	Flashing green	Standby module is starting to take over
Not lit	Solid green	Switchover is complete

Standby Led Behavior

To be completed.

Logging Messages

Log File

The log file displays messages about the activities and status of the management modules. Enter this command to display the messages:

Syntax: show logging

Displays log events.

For more information on command options available with the **show logging** command, see “CLI: Displaying the Event Log” in the “Troubleshooting” chapter of this guide.

An example of the log file listing is shown in Figure 15-31.

```
ProCurve(config)# show logging
Keys:   W=Warning   I=Information
        M=Major     D=Debug
----  Event Log listing:  Events Since Boot  ----
M 01/26/14 17:34:07 sys: 'System reboot due to Power Failure'
I 01/26/14 17:34:07 00061 system: -----
I 01/26/14 17:34:07 00062 system: Mgmt Module 2 went down without saving crash
                               information
I 01/26/14 17:36:14 00264 system: Mgmt Module 1 Failed Selftest
I 01/26/14 17:36:19 00068 chassis: Fabric 1 Inserted
I 01/26/14 17:36:19 00068 chassis: Fabric 2 Inserted
I 01/26/14 17:36:19 00068 chassis: Slot D Inserted
I 01/26/14 17:36:19 00690 udpf: DHCP relay agent feature enabled
I 01/26/14 17:36:19 00400 stack: Stack Protocol disabled
I 01/26/14 17:36:19 00128 tftp: Enable succeeded
I 01/26/14 17:36:19 00417 cdp: CDP enabled
I 01/26/14 17:36:19 00688 lldp: LLDP - enabled
I 01/26/14 17:36:19 00066 system: Mgmt Module 2 Booted
I 01/26/14 17:36:19 00260 system: Mgmt Module 2 Active
I 01/26/14 17:36:19 00066 system: Mgmt Module 1 Booted
I 01/26/14 17:36:19 00261 system: Mgmt Module 1 in Standby Mode
I 01/26/14 17:36:27 00375 chassis: Slot D Downloading
I 01/26/14 17:36:29 00376 chassis: Slot D Download Complete
I 01/26/14 17:36:44 00422 chassis: Slot D Ready
I 01/26/14 17:39:28 00179 mgr: SME CONSOLE Session - MANAGER Mode
I 01/26/14 21:49:10 00261 system: Mgmt Module 1 in Standby Mode
----  Bottom of Log : Events Listed = 21  ----
```

Figure 15-31. Log File Listing

Crash Files

Crash logs for all modules are always available on the active management module. The **copy crash-log** and **copy crash-data** commands can be used to copy the information to a file of your choice.

Syntax: copy crash-log [<slot-id> | mm] tftp <ip-address> <filename>

*Copies both the active and standby management modules' crash logs to a user-specified file. If no parameter is specified, files from **all** modules (management and interface) are concatenated.*

slot-id: *retrieves the crash log from the module in the specified slot.*

mm: *retrieves the crash logs from both management modules and concatenates them.*

Syntax: copy crash-data [<slot-id> | mm] tftp <ip-address> <filename>

*Copies both the active and standby management modules' crash data to a user-specified file. If no parameter is specified, files from **all** modules (management and interface) are concatenated.*

slot-id: *retrieves the crash data from the module in the specified slot.*

mm: *retrieves the crash data from both management modules and concatenates them.*

Displaying Saved Crash Information

You can display the saved crash information for each management module by using this command:

Syntax: show boot-history

Displays the system boot log.

An example of the output is shown in Figure 15-32.


```
ProCurve Switch 8200zl$ show boot-history

Mgmt Module 1 -- Saved Crash Information (most recent first):
=====
Mgmt Module 1 in Active Mode went down:  11/07/05 14:48:36
Operator warm reload from CONSOLE session.

Mgmt Module 1 in Active Mode went down:  11/07/05 11:43:10
Operator cold reboot from CONSOLE session.

Mgmt Module 2 -- Saved Crash Information (most recent first):
=====
  No Saved Crash Information
```

Figure 15-32. An Example of the System Boot Log File

Notes on How the Active Module is Determined

Both management modules run selftest routines to determine which module becomes the active management module and which becomes the standby management module. The module that was last active in the chassis is given precedence and becomes the “active” module. This module will be the one that is booted going forward. If a module fails selftest and is unable to communicate with the other module, it does not take control as the management module. The other management module will take control and become the active module.

If both modules fail selftest, the fault LED flashes and neither module is operational.

Note

You are not allowed to switchover to a management module that is not in standby mode. The module must have passed selftest and be in standby mode.

The entire boot decision process works as follows:

1. If there is only one management module, that is the active management module.
2. If one module is already booted and operational, a newly inserted module or the other management module booting will always become the standby module. The standby module does not become active unless a switchover occurs.
3. If there are two management modules and one fails selftest, the one that passes selftest becomes the active management module.
4. If only one of two modules was ever booted in the chassis, that module is given precedence.
5. The module that was active on the last boot becomes the active management module. This guarantees that the active module has the latest configuration data.
6. If both management modules have previously booted in this chassis and were “active” the last time booted, the module that booted most recently becomes the active management module.
7. If none of the above conditions are applicable, the module in the lowest slot becomes the active management module.

Diagram of Decision Process

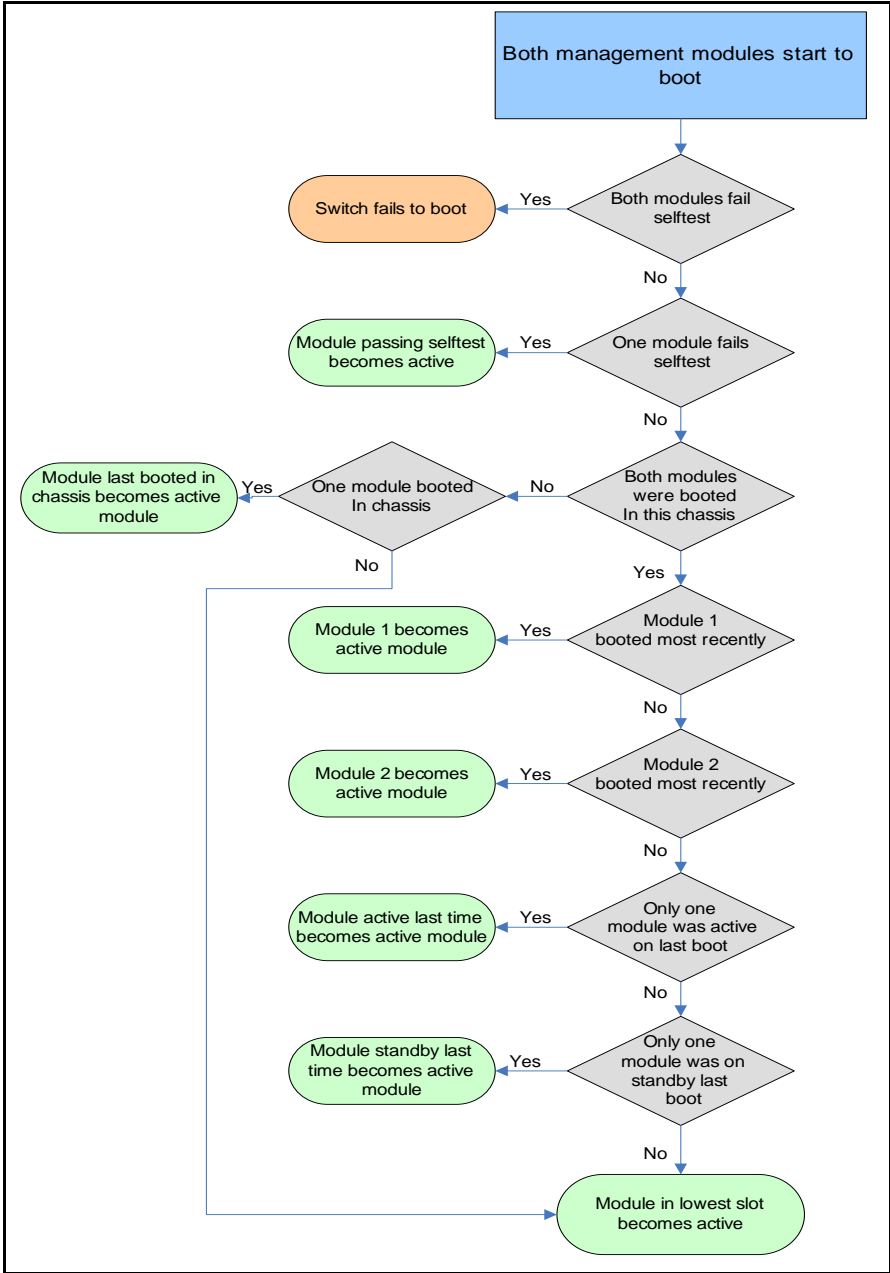


Figure 15-33. Active Module Decision Flow Chart at Boot

Event Log Messages

#	System Message	Severity	Description
1	Mgmt module [1 or 2] went down without saving crash information	info	The specified management module went down without saving the crash information. RMON_BOOT_NO_CRASH_RECORD
2	Mgmt module [1 or 2] went down	info	The specified management module was rebooted. RMON_BOOT_CRASH_RECORD0
4	Boot-up selftest failed	fatal	The boot up selftest of the management module failed. This message also appears for an interface module that fails selftest. RMON_BOOT_SELFTEST_FAILURE
5	System booted or Mgmt Module [1 or 2] Booted	info	This message appears at the end of the boot process. RMON_BOOT_COMPLETE
8	Mgmt Module [1 or 2] Active	info	The specified management module is the active management module RMON_SYSTEM_MGMT_MOD_ACTIVE
9	Mgmt Module [1 or 2] in Standby Mode	info	The specified management module is in standby mode. RMON_SYSTEM_MGMT_MOD_STANDBY
11	Mgmt Module [1 or 2] Offline (redundancy disabled)	info	The no redundancy management-module command was issued and the specified management module went offline. RMON_SYSTEM_MGMT_OFFLINE
12	Mgmt Module [1 or 2] Failed Selftest	warn	The specified management module failed selftest and will not become an active or standby module. RMON_SYSTEM_MGMT_FAILED
13	Lost Communication with Mgmt Module [1 or 2]	warn	A management module failed to receive heartbeats from the other management module. RMON_SYSTEM_MGMT_LOST_COMM
14	Resetting Mgmt Module [1 or 2]	info	The specified management module is being reset. This can occur if the MM Reset button is pressed. RMON_SYSTEM_MGMT_RESET

#	System Message	Severity	Description
15	Mgmt Module [1 or 2] - Running different version of SW	info	The specified management module is running a different version of software from the other management module. RMON_SYSTEM_MGMT_OS_DIFF
16	Mgmt Module [1 or 2] - Failover occurred	warn	Switchover occurred. The specified management module is the newly active management module RMON_SYSTEM_MGMT_FAILOVER
17	Mgmt Module [1 or 2] - User initiated switchover occurred	info	The user has initiated switchover using the redundancy switchover command so that the other management module can become the active management module. RMON_SYSTEM_MGMT_SWITCHOVER
18	Mgmt Module [1 or 2] - Offline (incompatible SW versions)	warn	The software version in the specified management module is not compatible with the other module and it has gone offline. RMON_SYSTEM_MGMT_INCOMPAT_OS
19	Other management module is not in standby, shutdown request ignored	warn	A shutdown request is ignored because the standby module is not in standby mode. A management module must be in active or standby mode to be shut down. The module goes into a "down" state which allows you to safely swap it out. RMON_SYSTEM_MGMT_HSBUTTONERR
20	Mgmt Module [1 or 2] Offline (shutdown)	info	The specified management module is offline because of a shutdown. RMON_SYSTEM_MGMT_SHUTDOWN
22	Syncing [primary secondary] OS to standby	info	This message is logged when the OS begins synchronizing to the standby module. RMON_SYSTEM_OS_SYNC
23	Standby boot image updated, rebooting standby	info	This message is logged when the standby's management module boot image is overwritten, requiring a reboot. RMON_SYSTEM_STANDBY_REBOOT

Redundancy (Switches 8200zl)

Event Log Messages

#	System Message	Severity	Description
24	Initial active to standby sync started	info	Indicates the beginning of the initial synchronization of the active management module's flash image to the standby management module. RMON_SYSTEM_SYNC_BEGIN
25	Initial active to standby sync complete	info	Indicates the end of the initial synchronization of the active management module's flash image to the standby management module. RMON_SYSTEM_SYNC_END

File Transfers

Contents

Overview	A-3
Downloading Switch Software	A-3
General Software Download Rules	A-4
Using TFTP To Download Software from a Server	A-4
Menu: TFTP Download from a Server to Primary Flash	A-5
CLI: TFTP Download from a Server to Flash	A-7
Enabling TFTP	A-9
Using Auto-TFTP	A-10
Using Secure Copy and SFTP	A-11
How It Works	A-12
The SCP/SFTP Process	A-12
Disable TFTP and Auto-TFTP for Enhanced Security	A-13
Command Options	A-14
Authentication	A-15
SCP/SFTP Operating Notes	A-15
Troubleshooting SSH, SFTP, and SCP Operations	A-17
Using Xmodem to Download Switch Software From a PC or UNIX Workstation	A-19
Menu: Xmodem Download to Primary Flash	A-19
CLI: Xmodem Download from a PC or UNIX Workstation to Primary or Secondary Flash	A-20
Using USB to Transfer Files to and from the Switch	A-21
Using USB to Download Switch Software	A-22
Switch-to-Switch Download	A-23
Menu: Switch-to-Switch Download to Primary Flash	A-24
CLI: Switch-To-Switch Downloads	A-25
Using PCM+ to Update Switch Software	A-26
Copying Software Images	A-27

TFTP: Copying a Software Image to a Remote Host	A-27
Xmodem: Copying a Software Image from the Switch to a Serially Connected PC or UNIX Workstation	A-27
USB: Copying a Software Image to a USB Device	A-28
Transferring Switch Configurations	A-28
TFTP: Copying a Configuration File to a Remote Host	A-29
TFTP: Copying a Configuration File from a Remote Host	A-30
TFTP: Copying a Customized Command File to a Switch	A-30
Xmodem: Copying a Configuration File to a Serially Connected PC or UNIX Workstation	A-32
Xmodem: Copying a Configuration File from a Serially Connected PC or UNIX Workstation	A-32
USB: Copying a Configuration File to a USB Device	A-34
USB: Copying a Configuration File from a USB Device	A-34
Transferring ACL Command Files	A-35
TFTP: Uploading an ACL Command File from a TFTP Server	A-35
Xmodem: Uploading an ACL Command File from a Serially Connected PC or UNIX Workstation	A-37
USB: Uploading an ACL Command File from a USB Device . .	A-37
Copying Diagnostic Data to a Remote Host, USB Device, PC or UNIX Workstation	A-38
Copying Command Output to a Destination Device	A-39
Copying Event Log Output to a Destination Device	A-40
Copying Crash Data Content to a Destination Device	A-40
Copying Crash Log Data Content to a Destination Device	A-42
Enabling or Disabling the USB Port	A-44
Behavior of Autorun When USB Port is Disabled	A-45
Software Versions K.13.XX Operation	A-45
Software Version K.14.XX Operation	A-45
Using USB Autorun	A-46
How It Works	A-46
Security Considerations	A-47
Troubleshooting Autorun Operations	A-48
Configuring Autorun on the Switch	A-49
Enabling Secure Mode	A-49

Operating Notes and Restrictions A-50
Autorun and Configuring Passwords A-50
Viewing Autorun Configuration Information A-51

Overview

The switches covered in this guide support several methods for transferring files to and from a physically connected device, or via the network, including TFTP, Xmodem, and USB. This appendix explains how to download new switch software, upload or download switch configuration files and software images, and upload command files for configuring Access Control Lists (ACLs). It contains the following information:

- Downloading switch software (begins on this page)
- Copying software images (page A-28)
- Transferring switch configurations (begins on page A-29)
- Uploading ACL command files (begins on page A-36)
- Copying diagnostic data (begins on page A-39)
- Using USB Autorun (begins on page A-47)

Downloading Switch Software

ProCurve periodically provides switch software updates through the ProCurve Networking web site. For more information, refer to the support and warranty booklet shipped with the switch, or visit www.procurve.com and click on **software updates**. After you acquire a new software version, you can use one of the following methods for downloading software to the switch:

Software Download Feature	Default	Menu	CLI	Web
TFTP	n/a	page A-6	page A-8	—
Xmodem	n/a	page A-20	page A-21	—
USB	n/a	n/a	page A-22	—
Switch-to-Switch	n/a	page A-25	page A-26	—
Software Update Manager in PCM+	Refer to the documentation provided with PCM+.			

Note

This manual uses the terms *switch software* and *software image* to refer to the downloadable software files the switch uses to operate its networking features. Other terms sometimes include *Operating System*, or *OS*.

General Software Download Rules

- Switch software that you download via the menu interface always goes to primary flash.
- After a software download, you must reboot the switch to implement the new software. Until a reboot occurs, the switch continues to run on the software it was using before the download commenced.

Note

Downloading new switch software does not change the current switch configuration. The switch configuration is contained in separate files that can also be transferred. Refer to “Transferring Switch Configurations” on page A-28.

In most cases, if a power failure or other cause interrupts a flash image download, the switch reboots with the image previously stored in primary flash. In the unlikely event that the primary image is corrupted (which may occur if a download is interrupted by a power failure), the switch goes into boot ROM mode. In this case, use the boot ROM console to download a new image to primary flash. Refer to “Restoring a Flash Image” on page C-84.

Using TFTP To Download Software from a Server

This procedure assumes that:

- A software version for the switch has been stored on a TFTP server accessible to the switch. (The software file is typically available from the ProCurve Networking web site at www.procurve.com.)
- The switch is properly connected to your network and has already been configured with a compatible IP address and subnet mask.
- The TFTP server is accessible to the switch via IP.

Before you use the procedure, do the following:

- Obtain the IP address of the TFTP server in which the software file has been stored.
- If VLANs are configured on the switch, determine the name of the VLAN in which the TFTP server is operating.
- Determine the name of the software file stored in the TFTP server for the switch (for example, E0820.swi).

Note

If your TFTP server is a UNIX workstation, *ensure that the case (upper or lower) that you specify for the filename is the same case as the characters in the software filenames on the server.*

Menu: TFTP Download from a Server to Primary Flash

Note that the menu interface accesses only the primary flash.

1. In the console Main Menu, select **Download OS** to display the screen in figure A-1. (The term “OS”, or “operating system” refers to the switch software):

```
----- CONSOLE - MANAGER MODE -----
                          Download OS

Current Firmware revision : K.11.00

Method [TFTP] : TFTP
TFTP Server :

Remote File Name :

Actions->  Cancel    Edit    eXecute    Help

Select the file transfer method (TFTP and XMODEM are currently supported).
Use arrow keys to change field selection, <Space> to toggle field choices,
and <Enter> to go to Actions.
```

Figure A-1. Example of a Download OS (Software) Screen (Default Values)

2. Press [E] (for **E**dit).
3. Ensure that the **Method** field is set to **TFTP** (the default).
4. In the **TFTP Server** field, type in the IP address of the TFTP server in which the software file has been stored.
5. In the **Remote File Name** field, type the name of the software file. If you are using a UNIX system, remember that the filename is case-sensitive.
6. Press [Enter], then [X] (for **eX**ecute) to begin the software download. The following screen then appears:

```
----- CONSOLE - MANAGER MODE -----
                          Download OS

Current Firmware revision : E.08.00
Method [TFTP] : TFTP
TFTP Server : 10.28.227.105

Remote File Name : K.11.00.swi

                          Received 370,000 bytes of OS download.
+-----+
|*****|
+-----+
```

Figure A-2. Example of the Download OS (Software) Screen During a Download

A “progress” bar indicates the progress of the download. When the entire software file has been received, all activity on the switch halts and you will see **Validating and writing system software to FLASH...**

7. After the primary flash memory has been updated with the new software, you must reboot the switch to implement the newly downloaded software. Return to the Main Menu and press [6] (for **Reboot Switch**). You will then see this prompt:

```
Continue reboot of system? : No
```

Press the space bar once to change No to Yes, then press [Enter] to begin the reboot.

Note

When you use the menu interface to download a switch software, the new image is always stored in primary flash. Also, using the Reboot Switch command in the Main Menu always reboots the switch from primary flash. Rebooting the switch from the CLI gives you more options. Refer to “Rebooting the Switch” on page 6-19.

8. After you reboot the switch, confirm that the software downloaded correctly:
 - a. From the Main Menu, select **1. Status and Counters**, and from the Status and Counters menu, select **1. General System Information**
 - b. Check the **Firmware revision** line.

Troubleshooting TFTP Download Failures. When using the menu interface, if a TFTP download fails, the Download OS (Operating System, or software) screen indicates the failure.

Message Indicating
cause of TFTP Download
Failure

```
----- CONSOLE - MANAGER MODE -----  
Download OS  
  
Current Firmware revision : K.11.00  
  
Method [TFTP] : TFTP  
TFTP Server : 10.29.227.105  
  
Remote File Name : os  
  
Received 0 bytes of OS download.  
+-----+  
|                                             |  
+-----+  
  
Connection to 10.29.227.105 failed  
  
Press any key to continue
```

Figure A-3. Example of Message for Download Failure

To find more information on the cause of a download failure, examine the messages in the switch's Event Log by executing the **show log tftp** command from the CLI. Also:

- For more on the Event Log, see “Using the Event Log for Troubleshooting Switch Problems” on page C-27.
- For descriptions of individual Event Log messages, refer to the latest version of the *Event Log Message Reference Guide* for your switch, available on the ProCurve website. (See also “Getting Documentation From the Web” on page 1-6.)

Some of the causes of download failures include:

- Incorrect or unreachable address specified for the **TFTP Server** parameter. This may include network problems.
- Incorrect VLAN.
- Incorrect name specified for the **Remote File Name** parameter, or the specified file cannot be found on the TFTP server. This can also occur if the TFTP server is a UNIX machine and the case (upper or lower) for the filename on the server does not match the case for the filename entered for the **Remote File Name** parameter in the **Download OS** (Operating System, or software) screen.
- One or more of the switch's IP configuration parameters are incorrect.
- For a UNIX TFTP server, the file permissions for the software file do not allow the file to be copied.
- Another console session (through either a direct connection to a terminal device or through Telnet) was already running when you started the session in which the download was attempted.

Note

If an error occurs in which normal switch operation cannot be restored, the switch automatically reboots itself. In this case, an appropriate message is displayed after the switch reboots.

CLI: TFTP Download from a Server to Flash

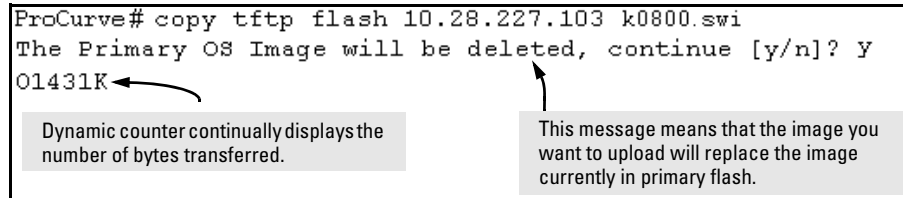
Syntax: copy tftp flash <ip-address> <remote-file> [< primary | secondary >]
[oobm]

This command automatically downloads a switch software file to primary or secondary flash. Note that if you do not specify the flash destination, the TFTP download defaults to primary flash.

For example, to download a switch software file named k0800.swi from a TFTP server with the IP address of 10.28.227.103 to primary flash:

1. Execute **copy** as shown below:

```
ProCurve# copy tftp flash 10.28.227.103 k0800.swi
The Primary OS Image will be deleted, continue [y/n]? Y
01431K
```



Dynamic counter continually displays the number of bytes transferred.

This message means that the image you want to upload will replace the image currently in primary flash.

Figure A-4. Example of the Command to Download an OS (Switch Software)

2. When the switch finishes downloading the software file from the server, it displays this progress message:

Validating and Writing System Software to FLASH ...

3. When the download finishes, you must reboot the switch to implement the newly downloaded software image. To do so, use one of the following commands:

Syntax: boot system flash < primary | secondary >

Boots from the selected flash.

Syntax: reload

Boots from the flash image and startup-config file. A switch covered in this guide (with multiple configuration files), also uses the current startup-config file.

(For more on these commands, refer to “Rebooting the Switch” on page 6-19.)

4. To confirm that the software downloaded correctly, execute **show system** and check the Firmware revision line.

For information on primary/secondary flash memory and the boot commands, refer to “Using Primary and Secondary Flash Image Options” on page 6-14.

Note

If you use **auto-tftp** to download a new image in a redundant management system, the active management module downloads the new image to both the active and standby modules. Rebooting after the auto-tftp process completes reboots the entire system.

Enabling TFTP

TFTP is enabled by default on the switch. If TFTP operation has been disabled, you can re-enable it by specifying TFTP client or server functionality with the **tftp <client | server>** command at the global configuration level.

Syntax: [no] tftp <client | server [listen <oobm|data|both>]>

Disables/re-enables TFTP for client or server functionality so that the switch can:

- *Use TFTP client functionality to access TFTP servers in the network to receive downloaded files.*
- *Use TFTP server functionality to upload files to other devices on the network.*
- *For switches that have a separate out-of-band management port, the **listen** parameter in a **server** configuration allows you to specify whether transfers take place through the out-of-band management (oobm) interface, the data interface, or both. Refer to Appendix I, “Networked Out-of-Band Management (OOBM)” in this guide for more information on out-of-band management.*

Usage Notes

To disable all TFTP client or server operation on the switch except for the auto-TFTP feature, enter the **no tftp <client | server>** command.

When ip ssh file transfer is used to enable SCP and SFTP functionality on the switch, this will disable TFTP client and server functionality. Once ip ssh file transfer is enabled, TFTP and auto-TFTP cannot be re-enabled from the CLI.

When TFTP is disabled, instances of TFTP in the CLI **copy** command and the Menu interface “Download OS” screen become unavailable.

The **no tftp <client | server>** command does not disable auto-TFTP operation. To disable an auto-TFTP command configured on the switch, use the **no auto-tftp** command described on page A-11 to remove the command entry from the switch’s configuration.

For information on how to configure TFTP file transfers on an IPv6 network, refer to the “IPv6 Management Features” chapter in the *IPv6 Configuration Guide* for your switch.

Using Auto-TFTP

The **auto-tftp** command allows you to configure the switch to download software automatically from a TFTP server.

How It Works. At switch startup, the auto-TFTP feature automatically downloads a specified software image to the switch from a specified TFTP server, then reboots the switch. To implement the process, you must first reboot the switch using one of the following methods:

- enter the **boot system flash primary** command in the CLI
- with the default flash boot image set to primary flash (the default), enter the **boot** or the **reload** command, or cycle the power to the switch. (To reset the boot image to primary flash, use **boot set-default flash primary**.)

Syntax: auto-tftp <ip-addr> <filename>

By default, auto-TFTP is disabled. This command configures the switch to automatically download the specified software file from the TFTP server at the specified IP address. The file is downloaded into primary flash memory at switch startup. The switch then automatically reboots from primary flash.

Notes: To enable auto-TFTP to copy a software image to primary flash memory, the version number of the downloaded software file (for example, K_14_01.swi) must be different from the version number currently in the primary flash image.

The current TFTP client status (enabled or disabled) does not affect auto-TFTP operation. (Refer to “Enabling TFTP” on page A-10.)

Completion of the auto-TFTP process may require several minutes while the switch executes the TFTP transfer to primary flash, and then reboots again.

*The **no** form of the command disables auto-TFTP operation by deleting the **auto-tftp** entry from the startup configuration. The **no auto-tftp** command does not affect the current TFTP-enabled configuration on the switch. However, entering the **ip ssh filetransfer** command automatically disables both **auto-tftp** and **tftp** operation.*

Using Secure Copy and SFTP

For some situations you may want to use a secure method to issue commands or copy files to the switch. By opening a secure, encrypted SSH session and enabling `ip ssh file transfer`, you can then use a third-party software application to take advantage of Secure Copy (SCP) and Secure ftp (SFTP). SCP and SFTP provide a secure alternative to TFTP for transferring information that may be sensitive (like switch configuration files) to and from the switch. Essentially you are creating a secure SSH tunnel as a way to transfer files with SFTP and SCP channels.

To use these commands you must install on the administrator workstation a third-party application software client that supports the SFTP and/or SCP functions. Some examples of software that supports SFTP and SCP are PuTTY, Open SSH, WinSCP, and SSH Secure Shell. Most of these are freeware and may be downloaded without cost or licensing from the internet. There are differences in the way these clients work, so be sure you also download the documentation.

As described earlier in this chapter you can use a TFTP client on the administrator workstation to update software images. This is a plain text mechanism and it connects to a standalone TFTP server or another ProCurve switch acting as a TFTP server to obtain the software image file(s). Using SCP and SFTP allows you to maintain your switches with greater security. You can also roll out new software images with automated scripts that make it easier to upgrade multiple switches simultaneously and securely.

SFTP (secure file transfer protocol) is unrelated to FTP, although there are some functional similarities. Once you set up an SFTP session through an SSH tunnel, some of the commands are the same as FTP commands. Certain commands are not allowed by the SFTP server on the switch, such as those that create files or folders. If you try to issue commands such as **create** or **remove** using SFTP the switch server returns an error message.

You can use SFTP just as you would TFTP to transfer files to and from the switch, but with SFTP your file transfers are encrypted and require authentication, so they are more secure than they would be using TFTP. SFTP works only with SSH version 2 (SSH v2).

Note

SFTP over SSH version 1 (SSH v1) is not supported. A request from either the client or the switch (or both) using SSH v1 generates an error message. The actual text of the error message differs, depending on the client software in use. Some examples are:

```
Protocol major versions differ: 2 vs. 1
Connection closed

Protocol major versions differ: 1 vs. 2
Connection closed

Received disconnect from <ip-addr>: /usr/local/
libexec/sftp-server: command not supported
Connection closed
```

SCP (secure copy) is an implementation of the BSD **rcp** (Berkeley UNIX remote copy) command tunneled through an SSH connection.

SCP is used to copy files to and from the switch when security is required. SCP works with both SSH v1 and SSH v2. Be aware that the most third-party software application clients that support SCP use SSHv1.

How It Works

The general process for using SCP and SFTP involves three steps:

1. Open an SSH tunnel between your computer and the switch if you haven't already done so. (This step assumes that you have already set up SSH on the switch.)
2. Execute **ip ssh filetransfer** to enable secure file transfer.
3. Use a third-party client application for SCP and SFTP commands.

The SCP/SFTP Process

To use SCP and SFTP:

1. Open an SSH session as you normally would to establish a secure encrypted tunnel between your computer and the switch. For more detailed directions on how to open an SSH session refer to the chapter titled "*Configuring Secure Shell (SSH)*" in the *Access Security Guide* for your switch. Please note that this is a one-time procedure for new switches or connections. If you have already done it once you should not need to do it a second time.
2. To enable secure file transfer on the switch (once you have an SSH session established between the switch and your computer), open a terminal window and type in the following command:

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

Disable TFTP and Auto-TFTP for Enhanced Security

Using the **ip ssh filetransfer** command to enable Secure FTP (SFTP) automatically disables TFTP and auto-TFTP (if either or both are enabled).

```
ProCurve(config)# ip ssh filetransfer
Tftp and auto-tftp have been disabled.
ProCurve(config)# sho run

Running configuration:

; J8697 Configuration Editor; Created on release #K.11.XX

hostname "ProCurve"
module 1 type J8702A
module 2 type J702A
vlan 1
    name "DEFAULT_VLAN"
    untagged A1-A24,B1-B24
    ip address 10.28.234.176 255.255.240.0
    exit
ip ssh filetransfer
no tftp-enable
password manager
password operator
```

Enabling SFTP automatically disables TFTP and auto-tftp and displays this message.

Viewing the configuration shows that SFTP is enabled and TFTP is disabled.

Figure A-5. Example of Switch Configuration with SFTP Enabled

If you enable SFTP, then later disable it, TFTP and auto-TFTP remain disabled unless they are explicitly re-enabled.

Operating rules are:

- The TFTP feature is enabled by default, and can be enabled or disabled through the CLI, the Menu interface, or an SNMP application. Auto-TFTP is disabled by default and must be configured through the CLI.

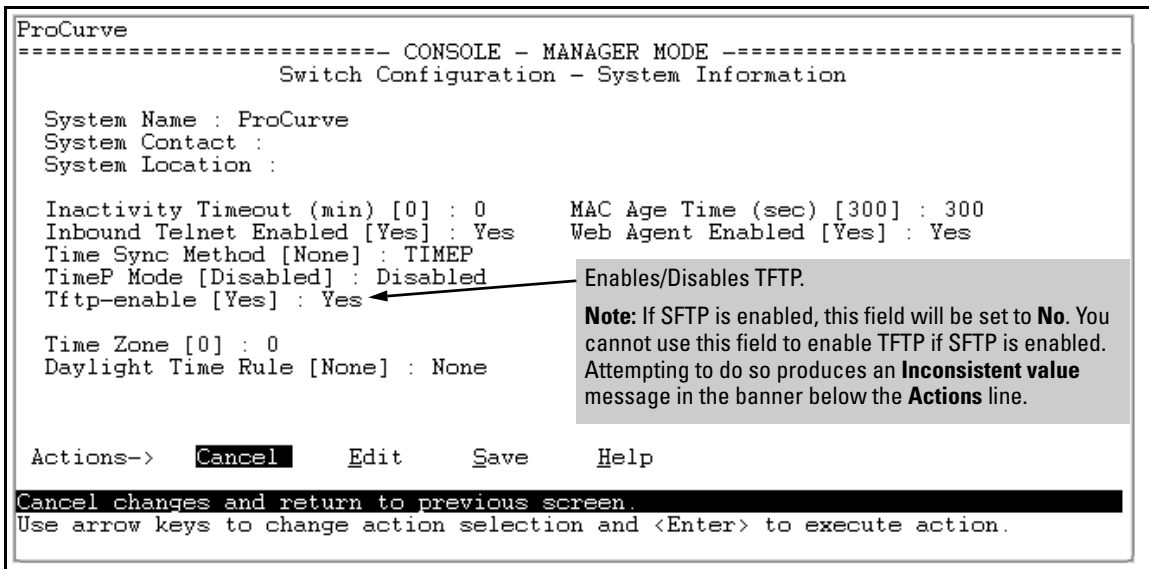


Figure A-6. Using the Menu Interface To Disable TFTP

- While SFTP is enabled, TFTP and auto-TFTP cannot be enabled from the CLI. Attempting to enable either non-secure TFTP option while SFTP is enabled produces one of the following messages in the CLI:

SFTP must be disabled before enabling tftp.

SFTP must be disabled before enabling auto-tftp.

Similarly, while SFTP is enabled, TFTP cannot be enabled using an SNMP management application. Attempting to do so generates an “inconsistent value” message. (An SNMP management application cannot be used to enable or disable auto-TFTP.)

- To enable SFTP by using an SNMP management application, you must first disable TFTP and, if configured, auto-TFTP on the switch. You can use either an SNMP application or the CLI to disable TFTP, but must use the CLI to disable auto-TFTP. The following two CLI commands disable TFTP and auto-TFTP on the switch.

Command Options

If you need to enable SSH v2 (which is required for SFTP) enter this command:

```
ProCurve(config)# ip ssh version 2
```

Note

As a matter of policy, administrators should *not* enable the SSHv1-only or the SSHv1-or-v2 advertisement modes. SSHv1 is supported on only some legacy switches (such as the ProCurve Series 2500 switches).

To confirm that SSH is enabled type in the command

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

Once you have confirmed that you have enabled an SSH session (with the **show ip ssh** command), enter **ip ssh filetransfer** so that SCP and/or SFTP can run. You can then open your third-party software client application to begin using the SCP or SFTP commands to safely transfer files or issue commands to the switch.

If you need to disable secure file transfer:

```
ProCurve(config)# no ip ssh filetransfer
```

Authentication

Switch memory allows up to ten public keys. This means the authentication and encryption keys you use for your third-party client SCP/SFTP software can differ from the keys you use for the SSH session, even though both SCP and SFTP use a secure SSH tunnel.

Note

SSH authentication is mutually exclusive with RADIUS servers.

Some clients such as PSCP (PuTTY SCP) automatically compare switch host keys for you. Other clients require you to manually copy and paste keys to the **\$HOME/.ssh/known_hosts** file. Whatever SCP/SFTP software tool you use, after installing the client software you must verify that the switch host keys are available to the client.

Because the third-party software utilities you may use for SCP/SFTP vary, you should refer to the documentation provided with the utility you select before performing this process.

SCP/SFTP Operating Notes

- Any attempts to use SCP or SFTP without using **ip ssh filetransfer** will cause the SCP or SFTP session to fail. Depending on the client software in use, you will receive an error message on the originating console, for example:

```
IP file transfer not enabled on the switch
```

- When an SFTP client connects, the switch provides a file system displaying all of its available files and folders. No file or directory creation is permitted by the user. Files may only be uploaded or downloaded, according to the permissions mask. All of the necessary files the switch will need are already in place on the switch. You do not need to (nor can you create) new files.
- The switch supports one SFTP session or one SCP session at a time.
- All files have read-write permission. Several SFTP commands, such as `create` or `remove`, are not allowed and return an error message. The switch displays the following files:

```

/
+---cfg
|   running-config
|   startup-config
+---log
|   crash-data
|   crash-data-a
|   crash-data-b
|   crash-data-c
|   crash-data-d           8212zl only
|   crash-data-e           "           "
|   crash-data-f           "           "
|   crash-data-g           8212zl only
|   crash-data-h           "           "
|   crash-data-I           "           "
|   crash-data-J           "           "
|   crash-data-K           "           "
|   crash-data-L           "           "
|   crash-log
|   crash-log-a
|   crash-log-b
|   crash-log-c
|   crash-log-d           8212zl only
|   crash-log-e           "           "
|   crash-log-f           "           "
|   crash-log-g           8212zl only
|   crash-log-h           "           "
|   crash-log-I           "           "
|   crash-log-J           "           "
|   crash-log-K           "           "
|   crash-log-L           "           "
|   event log
+---os
|   primary
|   secondary
\---ssh
    +---mgr_keys
  
```

```
    |   authorized_keys
    \---oper_keys
        authorized_keys
\---core           (this directory is not available on the 8212zl)
|   mm1.cor        management module or management function
|   im_a.cor       interface module (chassis switches only)
|   im_b.cor       interface module (chassis switches only)
|   im_1.cor       interface module (chassis switches only)
|   port_1-24.cor  core-dump for ports 1-24 (stackable switches only)
|   port_25-48.cor core-dump for ports 25-48 (stackable switches only)
```

- When using SFTP to copy a software image onto the switch, the command return takes only a few seconds. However, this does not mean that the transfer is complete, because the switch requires additional time (typically more than one minute) to write the image to flash in the background. To verify the file transfer has been completed, you can use the **show flash** command or look for a confirmation message in the log as in the following example:

```
I 01/09/09 16:17:07 00150 update: Primary Image
updated.
```

Once you have configured your switch to enable secure file transfers with SCP and SFTP, files can be copied to or from the switch in a secure (encrypted) environment and TFTP is no longer necessary.

Troubleshooting SSH, SFTP, and SCP Operations

You can verify secure file transfer operations by checking the switch's event log, or by viewing the error messages sent by the switch that most SCP and SFTP clients will print out on their console.

Note

Messages that are sent by the switch to the client depend on the client software in use to display them on the user console.

Broken SSH Connection. If an ssh connection is broken at the wrong moment (for instance, the link goes away or spanning tree brings down the link), a fatal exception would occur on the switch. If this happens, the switch will gracefully exit the session and produce an event log message indicating the cause of failure. The following three examples show the error messages that may appear in the log depending on the type of session that is running (SSH, SCP, or SFTP).


```
ssh: read error Bad file number, session aborted I 01/01/90 00:06:11 00636 ssh: sftp session from ::ffff:10.0.12.35 W 01/01/90 00:06:26 00641 ssh: sftp read error Bad file number, session aborted I 01/01/90 00:09:54 00637 ssh: scp session from ::ffff:10.0.12.35 W 01/01/90 ssh: scp read error Bad file number, session aborted
```

Note

The Bad file number is from the system error value and may differ depending on the cause of the failure. In the third example, the device file to read was closed as the device read was about to occur.

Attempt to Start a Session During a Flash Write. If you attempt to start an SCP (or SFTP) session while a flash write is in progress, the switch will not allow the SCP or SFTP session to start. Depending on the client software in use, the following error message may appear on the client console:

```
Received disconnect from 10.0.12.31: 2: Flash access in progress  
lost connection
```

Failure to Exit from a Previous Session. This next example shows the error message that may appear on the client console if a new SCP (or SFTP) session is started from a client before the previous client session has been closed (the switch requires approximately ten seconds to timeout the previous session):

```
Received disconnect from 10.0.12.31: 2: Wait for previous session to complete  
lost connection
```

Attempt to Start a Second Session. The switch supports only one SFTP session or one SCP session at a time. If a second session is initiated (for example, an SFTP session is running and then an SCP session is attempted), then the following error message may appear on the client console:

```
Received disconnect from 10.0.12.31: 2: Other SCP/SFTP session running  
lost connection
```

Using Xmodem to Download Switch Software From a PC or UNIX Workstation

This procedure assumes that:

- The switch is connected via the Console RS-232 port to a PC operating as a terminal. (Refer to the *Installation and Getting Started Guide* you received with the switch for information on connecting a PC as a terminal and running the switch console interface.)
- The switch software is stored on a disk drive in the PC.
- The terminal emulator you are using includes the Xmodem binary transfer feature. (For example, in the HyperTerminal application included with Windows NT, you would use the **Send File** option in the **Transfer** dropdown menu.)

Menu: Xmodem Download to Primary Flash

Note that the menu interface accesses only the primary flash.

1. From the console Main Menu, select
7. Download OS
2. Press **[E]** (for **E**dit).
3. Use the Space bar to select **XMODEM** in the **Method** field.
4. Press **[Enter]**, then **[X]** (for **eX**ecute) to begin the software download. The following message then appears:

**Press enter and then initiate Xmodem transfer
from the attached computer....**

5. Press **[Enter]** and then execute the terminal emulator command(s) to begin Xmodem binary transfer. For example, using HyperTerminal:
 - a. Click on **Transfer**, then **Send File**.
 - b. Type the file path and name in the Filename field.
 - c. In the Protocol field, select **Xmodem**.
 - d. Click on the **[Send]** button.

The download will then commence. It can take several minutes, depending on the baud rate set in the switch and in your terminal emulator.

6. After the primary flash memory has been updated with the new software, you must reboot the switch to implement the newly downloaded software. Return to the Main Menu and press [6] (for **Reboot Switch**). You will then see the following prompt:

Continue reboot of system? : No

Press the space bar once to change No to Yes, then press [Enter] to begin the reboot.

7. To confirm that the software downloaded correctly:
 - a. From the Main Menu, select

1. Status and Counters

1. General System Information

- b. Check the **Firmware revision** line.

CLI: Xmodem Download from a PC or UNIX Workstation to Primary or Secondary Flash

Using Xmodem and a terminal emulator, you can download a software file to either primary or secondary flash.

Syntax: copy xmodem flash [< primary | secondary >]

Downloads a software file to primary or secondary flash. If you do not specify the flash destination, the Xmodem download defaults to primary flash.

For example, to download a switch software file named E0822.swi from a PC (running a terminal emulator program such as HyperTerminal) to primary flash:

1. Execute the following command in the CLI:

```
ProCurve# copy xmodem flash
The Primary OS Image will be deleted, continue [y/n]? y
Press 'Enter' and start XMODEM on your host...
```

2. Execute the terminal emulator commands to begin the Xmodem transfer. For example, using HyperTerminal:
 - a. Click on **Transfer**, then **Send File**.
 - b. Type the file path and name in the Filename field.

- c. In the Protocol field, select **Xmodem**.
- d. Click on the **[Send]** button.

The download can take several minutes, depending on the baud rate used in the transfer.

3. When the download finishes, you must reboot the switch to implement the newly downloaded software. To do so, use one of the following commands:

Syntax: boot system flash <primary | secondary>

Reboots from the selected flash.

Syntax: reload

Reboots from the flash image currently in use.

(For more on these commands, see “Rebooting the Switch” on page 6-19.)

4. To confirm that the software downloaded correctly:

```
ProCurve> show system
```

Check the **Firmware revision** line. It should show the software version that you downloaded in the preceding steps.

If you need information on primary/secondary flash memory and the boot commands, refer to “Using Primary and Secondary Flash Image Options” on page 6-14.

Using USB to Transfer Files to and from the Switch

The switch’s USB port (labeled as *Auxiliary Port*) allows the use of a USB flash drive for copying configuration files to and from the switch. Beginning with software release K_12_XX or later, **copy** commands that used either **tftp** or **xmodem**, now include an additional option for **usb** as a source or destination for file transfers.

Operating rules and restrictions on USB usage are:

- Unformatted USB flash drives must first be formatted on a PC (Windows FAT format). For devices with multiple partitions, only the first partition is supported. Devices with secure partitions are not supported.
- If they already exist on the device, sub-directories are supported. When specifying a <filename>, you must enter either the individual file name (if at the root) or the full path name (for example, /subdir/filename).

- To view the contents of a USB flash drive, use the **dir** command. This will list all files and directories at the root. To view the contents of a directory, you must specify the subdirectory name (that is, **dir <subdirectory>**).
- The USB port supports connection to a single USB device. USB hubs to add more ports are not supported.

Note

Some USB flash drives may not be supported on your switch. Consult the latest *Release Notes* for information on supported devices.

Using USB to Download Switch Software

This procedure assumes that:

- A software version for the switch has been stored on a USB flash drive. (The latest software file is typically available from the ProCurve Networking web site at www.procurve.com.)
- The USB device has been plugged into the switch's USB port.

Before you use the procedure:

- Determine the name of the software file stored on the USB flash drive (for example, k0800.swi).
- Decide whether the image will be installed in the primary or secondary flash. (For more on primary/secondary flash memory and related boot commands, refer to “Using Primary and Secondary Flash Image Options” on page 6-14.)

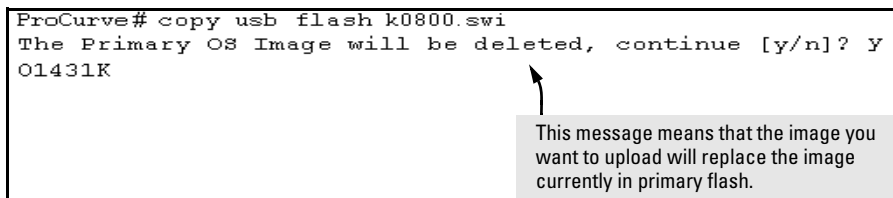
Syntax: copy usb flash <filename> [< primary | secondary >]

This command automatically downloads a switch software file to primary or secondary flash. Note that if you do not specify the flash destination, the USB download defaults to primary flash.

For example, to copy a switch software file named k0800.swi from a USB device to primary flash:

1. Execute **copy** as shown below:

```
ProCurve# copy usb flash k0800.swi
The Primary OS Image will be deleted, continue [y/n]? Y
01431K
```



This message means that the image you want to upload will replace the image currently in primary flash.

Figure A-7. Example of the Command to Copy Switch Software from USB

2. When the switch finishes copying the software file from the USB device, it displays this progress message:

Validating and Writing System Software to the Filesystem...

3. When the copy finishes, you must reboot the switch to implement the newly loaded software. To do so, use one of the following commands:

Syntax: boot system flash < primary | secondary >

Boots from the selected flash.

Syntax: reload

Boots from the flash image and startup-config file. A switch covered in this guide (with multiple configuration files), also uses the current startup-config file.

(For more on these commands, refer to “Rebooting the Switch” on page 6-19.)

4. To confirm that the software downloaded correctly, execute **show system** and check the Firmware revision line.

Switch-to-Switch Download

You can use TFTP to transfer a software image between two switches of the same series. The menu interface enables you to transfer primary-to-primary or secondary-to-primary. The CLI enables all combinations of flash location options.

Menu: Switch-to-Switch Download to Primary Flash

Using the menu interface, you can download a switch software file from either the primary or secondary flash of one switch to the primary flash of another switch of the same series.

1. From the switch console Main Menu in the switch to receive the download, select **7. Download OS** screen.
2. Ensure that the **Method** parameter is set to **TFTP** (the default).
3. In the **TFTP Server** field, enter the IP address of the remote switch containing the software file you want to download.
4. For the **Remote File Name**, enter one of the following:
 - To download the software in the primary flash of the source switch, type **“flash”** in lowercase characters.
 - To download the software in the secondary flash of the source switch, type **/os/secondary**.
5. Press **[Enter]**, then **[X]** (for **eXecute**) to begin the software download.
6. A “progress” bar indicates the progress of the download. When the entire switch software download has been received, all activity on the switch halts and the following messages appear:

Validating and writing system software to FLASH...

7. After the primary flash memory has been updated with the new software, you must reboot the switch to implement the newly downloaded software. Return to the Main Menu and press **[6]** (for **Reboot Switch**). You will then see this prompt:

Continue reboot of system? : No

Press the space bar once to change No to Yes, then press **[Enter]** to begin the reboot.

8. To confirm that the software downloaded correctly:
 - a. From the Main Menu, select

Status and Counters

General System Information

- b. Check the **Firmware revision** line.

CLI: Switch-To-Switch Downloads

Where two switches in your network belong to the same series, you can download a software image between them by initiating a **copy tftp** command from the destination switch. The options for this CLI feature include:

- Copy from primary flash in the source to either primary or secondary in the destination.
- Copy from either primary or secondary flash in the source to either primary or secondary flash in the destination.

Downloading from Primary Only.

Syntax: copy tftp flash < ip-addr > flash [primary | secondary] [oobm]

This command (executed in the destination switch) downloads the software flash in the source switch's primary flash to either the primary or secondary flash in the destination switch.

*For switches that have a separate out-of-band management port, the **oobm** parameter specifies that the TFTP traffic must come in through the out-of-band management interface. If this parameter is not specified, the TFTP traffic comes in through the data interface. The **oobm** parameter is not available on switches that do not have a separate out-of-band management port. Refer to Appendix I, "Network Out-of-Band Management" in this guide for more information on out-of-band management.*

If you do not specify either a primary or secondary flash location for the destination, the download automatically goes to primary flash.

For example, to download a software file from primary flash in a switch with an IP address of 10.29.227.103 to the primary flash in the destination switch, you would execute the following command in the destination switch's CLI:

```
ProCurve# copy tftp flash 10.29.227.103 flash
Device will be rebooted, do you want to continue [y/n] Y
00107K
```

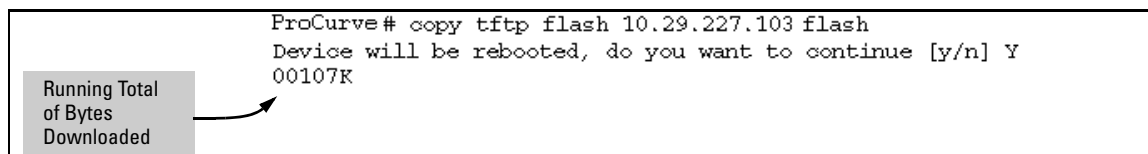


Figure A-8. Switch-To-Switch, from Primary in Source to Either Flash in Destination

Downloading from Either Flash in the Source Switch to Either Flash in the Destination Switch.

Syntax: copy tftp flash < ip-addr > < /os/primary > | < /os/secondary > [primary | secondary] [oobm]

This command (executed in the destination switch) gives you the most options for downloading between switches. If you do not specify either a primary or secondary flash location for the destination, the download automatically goes to primary flash.

*For switches that have a separate out-of-band management port, the **oobm** parameter specifies that the TFTP traffic must come in through the out-of-band management interface. If this parameter is not specified, the TFTP traffic comes in through the data interface. The **oobm** parameter is not available on switches that do not have a separate out-of-band management port. Refer to Appendix I, “Network Out-of-Band Management” in this guide for more information on out-of-band management.*

For example, to download a software file from secondary flash in a switch with an IP address of 10.28.227.103 to the secondary flash in a destination switch, you would execute the following command in the destination switch's CLI:

```
ProCurve# copy tftp flash 10.29.227.103 /os/secondary secondary
Device will be rebooted, do you want to continue [y/n] Y
01084K
```

Figure A-9. Switch-to-Switch, from Either Flash in Source to Either Flash in Destination

Using PCM+ to Update Switch Software

ProCurve Manager Plus includes a software update utility for updating on ProCurve switch products. For further information, refer to the *Getting Started Guide* and the *Administrator's Guide*, provided electronically with the application.

Copying Software Images

Using the CLI commands described in this section, you can copy software images from the switch to another device using tftp, xmodem, or usb.

Note

For details on how switch memory operates, including primary and secondary flash, refer to Chapter 6, “Switch Memory and Configuration”.

TFTP: Copying a Software Image to a Remote Host

Syntax: copy flash tftp < ip-addr > < filename > [oobm]

This command copies the primary flash image to a TFTP server.

*For switches that have a separate out-of-band management port, the **oobm** parameter specifies that the transfer will be through the out-of-band management interface. If this parameter is not specified, the transfer will be through the data interface. The **oobm** parameter is not available on switches that do not have a separate out-of-band management port. Refer to Appendix I, “Network Out-of-Band Management” in this guide for more information on out-of-band management.*

For example, to copy the primary flash to a TFTP server having an IP address of 10.28.227.105:

```
ProCurve# copy flash tftp 10.28.227.105 k0800.swi
```

where k0800.swi is the filename given to the flash image being copied.

Xmodem: Copying a Software Image from the Switch to a Serially Connected PC or UNIX Workstation

To use this method, the switch must be connected via the serial port to a PC or UNIX workstation.

Syntax: copy flash xmodem < pc | unix >

Uses Xmodem to copy a designated configuration file from the switch to a PC or Unix workstation.

For example, to copy the primary flash image to a serially connected PC:

1. Execute the following command:

```
Procurve# copy xmodem flash
```

Press 'Enter' and start XMODEM on your host...
2. After you see the above prompt, press **[Enter]**.
3. Execute the terminal emulator commands to begin the file transfer.

USB: Copying a Software Image to a USB Device

To use this method, a USB flash memory device must be connected to the switch's USB port.

Syntax: copy flash usb <filename>

Uses the USB port to copy the primary flash image from the switch to a USB flash memory device.

For example, to copy the primary image to a USB flash drive:

1. Insert a USB device into the switch's USB port.
2. Execute the following command:

```
Procurve# copy flash usb k0800.swi
```

where `k0800.swi` is the name given to the primary flash image that is copied from the switch to the USB device.

Transferring Switch Configurations

Transfer Features

Feature	Page
Use TFTP to copy from a remote host to a config file.	A-31
Use TFTP to copy a config file to a remote host.	A-33
Use Xmodem to copy a configuration from a serially connected host to a config file.	A-33
Use Xmodem to copy a config file to a serially connected host.	A-33
Use USB to copy a configuration from a USB device to a config file.	A-35
Use USB to copy a config file to a USB device.	A-35

Using the CLI commands described in this section, you can copy switch configurations to and from a switch, or copy a software image to configure or replace an ACL in the switch configuration.

Note

For greater security, you can perform all TFTP operations using SFTP as described in the section on *Using Secure Copy and SFTP* on page A-12.

The **include-credentials** command can also be used to save passwords, secret keys, and other security credentials in the running config file. For more information, see the section on “Saving Security Credentials in a Config File” in the *Access Security Guide* for your switch.

TFTP: Copying a Configuration File to a Remote Host

Syntax: copy < startup-config | running-config > tftp < ip-addr > < remote-file >
[pc | unix] [oobm]
copy config < filename > tftp < ip-addr > < remote-file > [pc | unix] [oobm]

This command can copy a designated config file in the switch to a TFTP server. For more on multiple configuration files, refer to “Multiple Configuration Files” on page 6-26.

*For switches that have a separate out-of-band management port, the **oobm** parameter specifies that the transfer will be through the out-of-band management interface. If this parameter is not specified, the transfer will be through the data interface. The **oobm** parameter is not available on switches that do not have a separate out-of-band management port. Refer to Appendix I, “Network Out-of-Band Management” in this guide for more information on out-of-band management.*

For example, to upload the current startup configuration to a file named **sw8200** in the configs directory on drive “d” in a TFTP server having an IP address of 10.28.227.105:

```
ProCurve# copy startup-config tftp 10.28.227.105  
d:\configs\sw8200
```

TFTP: Copying a Configuration File from a Remote Host

Syntax: copy tftp < startup-config | running-config > < ip-address > < remote-file >
[pc | unix] [oobm]
copy tftp config < filename > < ip-address > < remote-file > [pc | unix]
[oobm]

This command can copy a configuration from a remote host to a designated config file in the switch. For more on multiple configuration files, refer to “Multiple Configuration Files” on page 6-26.

(Refer to “Using Primary and Secondary Flash Image Options” on page 6-14 for more on flash image use.)

*For switches that have a separate out-of-band management port, the **oobm** parameter specifies that the transfer will be through the out-of-band management interface. If this parameter is not specified, the transfer will be through the data interface. The **oobm** parameter is not available on switches that do not have a separate out-of-band management port. Refer to Appendix I, “Network Out-of-Band Management” in this guide for more information on out-of-band management.*

For example, to download a configuration file named **sw8200** in the **configs** directory on drive “**d**” in a remote host having an IP address of 10.28.227.105:

```
ProCurve# copy tftp startup-config 10.28.227.105  
d:\configs\sw8200
```

TFTP: Copying a Customized Command File to a Switch

Using the **copy tftp** command with the **show-tech** option provides the ability to copy a customized command file to the switch. When the **show tech custom** command is executed, the commands in the custom file are executed instead of the hard-coded list of commands. If no custom file is found, the current hard-coded list is executed. This list contains commands to display data such as the image stamp, running configuration, boot history, port settings, and so on.

Syntax: copy tftp show-tech <ipv4 or ipv6 address> <filename> [oobm]

Copy a customized command file to the switch.

*For switches that have a separate out-of-band management port, the **oobm** parameter specifies that the transfer will be through the out-of-band management interface. If this parameter is not specified, the transfer will be through the data interface. The **oobm** parameter is not available on switches that do not have a separate out-of-band management port. Refer to Appendix I, “Network Out-of-Band Management” in this guide for more information on out-of-band management.*

```
ProCurve(config)# copy tftp show-tech 10.10.10.3 commandfile1
```

Figure A-10. Example of Using the copy tftp show-tech Command to Upload a Customized Command File

Syntax: show tech custom

Executes the commands found in a custom file instead of the hard-coded list.

Note: *Exit the global config mode (if needed) before executing **show tech** commands.*

You can include **show tech** commands in the custom file, with the exception of **show tech custom**. For example, you can include the command **show tech all**.

If no custom file is found, a message displays stating “No SHOW-TECH file found.”

```
ProCurve# show tech custom  
No SHOW-TECH file found.
```

No custom file was uploaded with the **copy tftp show-tech** command

Figure A-11. Example of the show tech custom Command

Xmodem: Copying a Configuration File to a Serially Connected PC or UNIX Workstation

To use this method, the switch must be connected via the serial port to a PC or UNIX workstation. You will need to:

- Determine a filename to use.
- Know the directory path you will use to store the configuration file.

Syntax: copy < startup-config | running-config > xmodem < pc | unix >
copy config < filename > xmodem < pc | unix >

Uses Xmodem to copy a designated configuration file from the switch to a PC or Unix workstation. For more on multiple configuration files, refer to “Multiple Configuration Files” on page 6-26.

For example, to copy a configuration file to a PC serially connected to the switch:

1. Determine the file name and directory location on the PC.
2. Execute the following command:

```
ProCurve# copy startup-config xmodem pc  
Press 'Enter' and start XMODEM on your host...
```

3. After you see the above prompt, press **[Enter]**.
4. Execute the terminal emulator commands to begin the file transfer.

Xmodem: Copying a Configuration File from a Serially Connected PC or UNIX Workstation

To use this method, the switch must be connected via the serial port to a PC or UNIX workstation on which is stored the configuration file you want to copy. To complete the copying, you will need to know the name of the file to copy and the drive and directory location of the file.

Syntax: copy xmodem startup-config < pc | unix >
copy xmodem config < filename > < pc | unix >

Copies a configuration file from a serially connected PC or UNIX workstation to a designated configuration file on the switch. For more on multiple configuration files, refer to “Multiple Configuration Files” on page 6-26.

For example, to copy a configuration file from a PC serially connected to the switch:

1. Execute the following command:

```
ProCurve# copy xmodem startup-config pc
Device will be rebooted, do you want to continue [y/n]? y
Press 'Enter' and start XMODEM on your host...
```

2. After you see the above prompt, press **[Enter]**.
3. Execute the terminal emulator commands to begin the file transfer.
4. When the download finishes, you must reboot the switch to implement the newly downloaded software. To do so, use one of the following commands:

Syntax: boot system flash [primary | secondary]
boot system flash [config < filename >

Switches boot from the designated configuration file. For more on multiple configuration files, refer to “Multiple Configuration Files” on page 6-26.

Syntax: reload

Reboots from the flash image currently in use.

(For more on these commands, refer to “Rebooting the Switch” on page 6-19.)

USB: Copying a Configuration File to a USB Device

To use this method, a USB flash memory device must be connected to the switch's USB port.

Syntax: copy startup-config usb < filename>
copy running-config usb < filename >

Uses the USB port to copy a designated configuration file from the switch to a USB flash memory device. For more on multiple configuration files, refer to “Multiple Configuration Files” on page 6-26.

For example, to copy the startup configuration file to a USB flash drive:

1. Insert a USB device into the switch's USB port.
2. Execute the following command:

```
Procurve# copy startup-config usb procurve-config
```

where `procurve-config` is the name given to the configuration file that is copied from the switch to the USB device.

USB: Copying a Configuration File from a USB Device

To use this method, the switch must be connected via the USB port to a USB flash drive on which is stored the configuration file you want to copy. To execute the command, you will need to know the name of the file to copy.

Syntax: copy usb startup-config < filename >

Copies a configuration file from a USB device to the startup configuration file on the switch.

For example, to copy a configuration file from a USB device to the switch:

1. Insert a USB device into the switch's USB port.
2. Execute the following command:

```
Procurve# copy usb startup-config procurve-config
```

where `procurve-config` is the name of the file to copy.

3. At the prompt, press **[Enter]** to reboot the switch and implement the newly downloaded software.

Transferring ACL Command Files

This section describes how to upload and execute a command file to the switch for configuring or replacing an Access Control List (ACL) in the switch configuration. Such files should contain only ACE (Access Control Entry) commands. For more on this general topic, including an example of an ACL command file created offline, refer to the section titled “Editing ACLs and Creating an ACL Offline” in the “Access Control Lists (ACLs)” chapter of the latest *Access Security Guide* for your switch.

TFTP: Uploading an ACL Command File from a TFTP Server

Syntax: `copy tftp command-file < ip-addr > < filename.txt > < unix | pc > [oobm]`

where:

`< ip-addr >` = *The IP address of a TFTP server available to the switch*

`< filename.txt >` = *A text file containing ACL commands and stored in the TFTP directory of the server identified by < ip-addr >*

`< unix | pc >` = *The type of workstation used for serial, Telnet, or SSH access to the switch CLI*

`[oobm]` = *For switches that have a separate out-of-band management port, specifies that the transfer will be through the out-of-band management interface. (Default is transfer through the data interface.)*

This command copies and executes the named text file from the specified TFTP server address and executes the ACL commands in the file. Depending on the ACL commands used, this action does one of the following in the running-config file:

- *Creates a new ACL.*
- *Replaces an existing ACL. (Refer to “Creating an ACL Offline” in the “Access Control Lists (ACLs)” chapter in the latest Access Security Guide for your switch.)*
- *Adds to an existing ACL.*

For example, suppose you:

1. Created an ACL command file named **vlan10_in.txt** to update an existing ACL.
2. Copied the file to a TFTP server at 18.38.124.16.

Using a PC workstation, you then execute the following from the CLI to upload the file to the switch and implement the ACL commands it contains:

```
ProCurve(config)# copy tftp command-file 18.38.124.16  
vlan10_in.txt pc
```

The switch displays this message:

```
Running configuration may change, do you want to continue  
[y/n]?
```

To continue with the upload, press the **[Y]** key. To abort the upload, press the **[N]** key. Note that if the switch detects an illegal (non-ACL) command in the file, it bypasses the illegal command, displays a notice as shown in figure A-12, and continues to implement the remaining ACL commands in the file.

```
ProCurve(config)# copy tftp command-file 10.38.124.16 vlan10_in.txt pc
Running configuration may change, do you want to continue [y/n]? y
 1. ip access-list extended "155"
 2. deny tcp 0.0.0.0 255.255.255.255 10.10.10.2 0.0.0.0 eq 23 log
 3. permit ip 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255
 4. show running
Command files are limited to access-list commands.
 5. exit
ProCurve(config)# show running

Running configuration:

; J8697A Configuration Editor; Created on release # K.11.00

hostname " ProCurve "
cdp run
module 1 type J8702A
ip default-gateway 10.38.248.1
logging 18.38.227.2
snmp-server community "public" Unrestricted
ip access-list extended "155"
deny tcp 0.0.0.0 255.255.255.255 10.10.10.2 0.0.0.0 eq 23 log
permit ip 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255
exit
-----
:
:
```

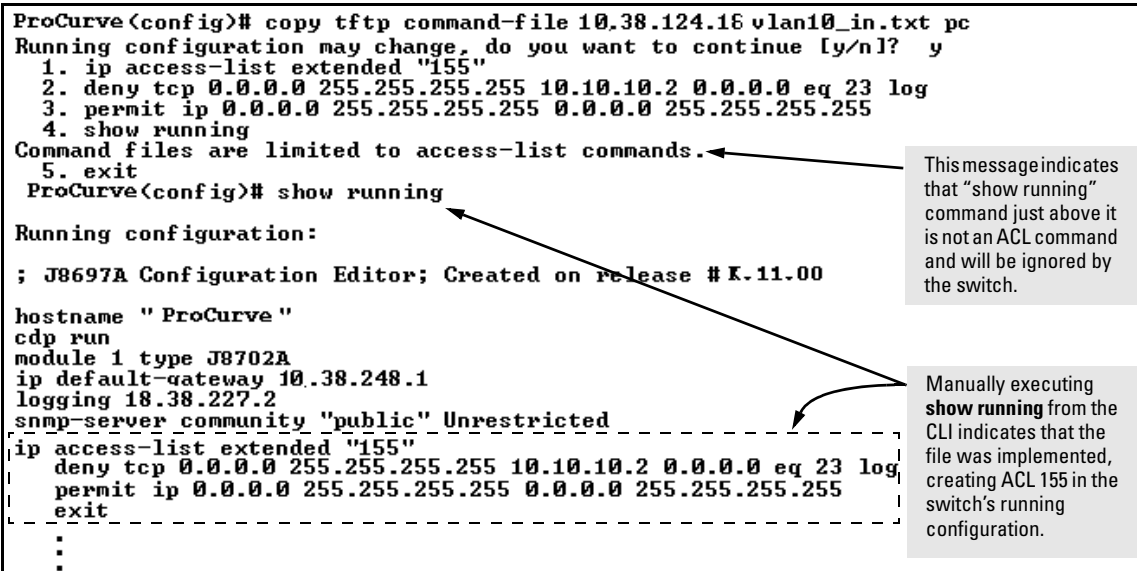


Figure A-12. Example of Using the Copy Command to Download and Configure an ACL

Xmodem: Uploading an ACL Command File from a Serially Connected PC or UNIX Workstation

Syntax: copy xmodem command-file < unix | pc >

Uses Xmodem to copy and executes an ACL command from a PC or Unix workstation. Depending on the ACL commands used, this action does one of the following in the running-config file:

- *Creates a new ACL.*
- *Replaces an existing ACL. (Refer to “Creating an ACL Offline” in the “Access Control Lists (ACLs)” chapter in the latest Access Security Guide for your switch.)*
- *Adds to an existing ACL.*

USB: Uploading an ACL Command File from a USB Device

Syntax: copy usb command-file < filename.txt > < unix | pc >

where:

< filename.txt > = A text file containing ACL commands and stored in the USB flash drive.

< unix | pc > = The type of workstation used to create the text file.

This command copies and executes the named text file from a USB flash drive and executes the ACL commands in the file. Depending on the ACL commands used, this action does one of the following in the running-config file:

- *Creates a new ACL.*
- *Replaces an existing ACL. (Refer to “Creating an ACL Offline” in the “Access Control Lists (ACLs)” chapter in the latest Access Security Guide for your switch.)*
- *Adds to an existing ACL.*

For example, suppose you:

1. Created an ACL command file named **vlan10_in.txt** to update an existing ACL.
2. Copied the file to a USB flash drive.

Using a PC workstation, you then execute the following from the CLI to upload the file to the switch and implement the ACL commands it contains:

```
ProCurve(config)# copy usb command-file vlan10_in.txt pc
```

The switch displays this message:

```
Running configuration may change, do you want to continue  
[y/n]?
```

To continue with the upload, press the **[Y]** key. To abort the upload, press the **[N]** key. Note that if the switch detects an illegal (non-ACL) command in the file, it bypasses the illegal command, displays a notice (as in the tftp example shown in Figure A-12 on page A-37), and continues to implement the remaining ACL commands in the file.

Copying Diagnostic Data to a Remote Host, USB Device, PC or UNIX Workstation

You can use the CLI to copy the following types of switch data to a text file in a destination device:

- **Command Output:** Sends the output of a switch CLI command as a file on the destination device.
- **Event Log:** Copies the switch's Event Log into a file on the destination device.
- **Crash Data:** software-specific data useful for determining the reason for a system crash.
- **Crash Log:** Processor-Specific operating data useful for determining the reason for a system crash.

The destination device and copy method options are as follows (CLI key word is in bold):

- Remote Host via **TFTP**.
- Physically connected USB flash drive via the switch's **USB** port.
- Serially connected PC or UNIX workstation via **Xmodem**.

Copying Command Output to a Destination Device

Syntax: `copy command-output < "cli-command" > tftp < ip-address > < filepath-filename > [oobm]`

`copy command-output < "cli-command" > usb < filename >`

`copy command-output < "cli-command" > xmodem`

These commands direct the displayed output of a CLI command to a remote host, attached USB device, or to a serially connected PC or UNIX workstation.

*For switches that have a separate out-of-band management port, the **oobm** parameter specifies that the transfer will be through the out-of-band management interface. If this parameter is not specified, the transfer will be through the data interface. The **oobm** parameter is not available on switches that do not have a separate out-of-band management port. Refer to Appendix I, "Network Out-of-Band Management" in this guide for more information on out-of-band management.*

For example, to use Xmodem to copy the output of **show config** to a serially connected PC:

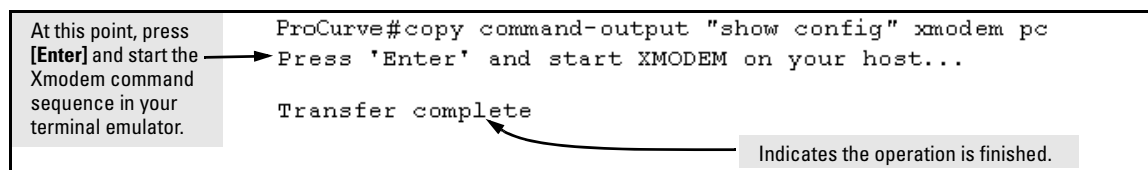


Figure A-13. Example of Sending Command Output to a File on an Attached PC

Note

The command you specify must be enclosed in double-quote marks.

Copying Event Log Output to a Destination Device

Syntax: `copy event-log tftp < ip-address > < filepath_filename > [oobm]`

`copy event-log usb < filename >`

`copy event-log xmodem <filename>`

These commands copy the Event Log content to a remote host, attached USB device, or to a serially connected PC or UNIX workstation.

*For switches that have a separate out-of-band management port, the **oobm** parameter specifies that the transfer will be through the out-of-band management interface. If this parameter is not specified, the transfer will be through the data interface. The **oobm** parameter is not available on switches that do not have a separate out-of-band management port. Refer to Appendix I, “Network Out-of-Band Management” in this guide for more information on out-of-band management.*

For example, to copy the event log to a PC connected to the switch:

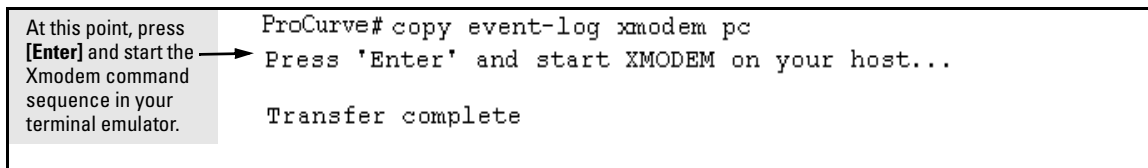


Figure A-14. Example of Sending Event Log Content to a File on an Attached PC

Copying Crash Data Content to a Destination Device

This command uses TFTP, USB, or Xmodem to copy the Crash Data content to a destination device. You can copy individual slot information or the management module’s switch information. If you do not specify either, the command defaults to the management function’s data.

Syntax: copy crash-data [<slot-id | master>] tftp <ip-address> <filename> [oobm]

copy crash-data [<slot-id | mm>] usb <filename>

copy crash-data [<slot-id | mm>] xmodem

where: slot-id a - h, and retrieves the crash log or crash data from the processor on the module in the specified slot.

mm Retrieves crash log or crash data from the switch's chassis processor. When "mm" is specified, crash files from both management modules are copied.

oobm For switches that have a separate out-of-band management port, specifies that the transfer will be through the out-of-band management interface. (Default is transfer through the data interface.)

These commands copy the crash data content to a remote host, attached USB device, or to a serially connected PC or UNIX workstation. You can copy individual slot information or the management module (mm) switch information. If you do not specify either, the command defaults to the mm data.

For example, to copy the switch's crash data to a file in a PC:

At this point, press [Enter] and start the Xmodem command sequence in your terminal emulator.	<pre> ProCurve(config)# copy crash-data xmodem pc Press 'Enter' and start XMODEM on your host... . Transfer complete </pre>
--	---

Figure A-15. Example of Copying Switch Crash Data Content to a PC

Copying Crash Data with Redundant Management. When you are using redundant management, the **copy crash-data** command operates somewhat differently.

Syntax: copy crash-data [<slot-id> | mm] tftp <ip-address> <filename> [oobm]

Copies both the active and standby management modules' crash data to a user-specified file. If no parameter is specified, files from all modules (management and interface) are concatenated.

slot-id: *retrieves the crash data from the module in the specified slot.*

mm: *retrieves the crash data from both management modules and concatenates them.*

oobm: *For switches that have a separate out-of-band management port, specifies that the transfer will be through the out-of-band management interface. (Default is transfer through the data interface.)*

Copying Crash Log Data Content to a Destination Device

Syntax: copy crash-log [<slot-id | mm>] tftp <ip-address>
<filepath and filename> [oobm]

copy crash-log [<slot-id | mm>] usb <filename>

copy crash-log [<slot-id | mm>] xmodem

where: slot-id = **a - h**, and retrieves the crash log from the processor on the module in the specified slot.

mm *Retrieves the crash log from the switch's chassis processor. When mm is specified, crash files from both management modules are copied.*

oobm *For switches that have a separate out-of-band management port, specifies that the transfer will be through the out-of-band management interface. (Default is transfer through the data interface.)*

These commands copy the Crash Log content to a remote host, attached USB device, or to a serially connected PC or UNIX workstation. You can copy individual slot information or the management module (mm) switch information. If you do not specify either, the command defaults to the mm data.

For example, to copy the Crash Log for slot C to a file in a PC connected to the switch:

File Transfers

Copying Diagnostic Data to a Remote Host, USB Device, PC or UNIX Workstation

At this point, press [Enter] and start the Xmodem command sequence in your terminal emulator.	<pre>ProCurve (config)# copy crash-log c xmodem Press 'Enter' and start XMODEM on your host... Transfer complete</pre>
--	--

Figure A-16. Example of sending a Crash Log for Slot C to a File on an Attached PC

Copying Crash Logs with Redundant Management. When you are using redundant management, the **copy crash-log** command operates somewhat differently.

Syntax: `copy crash-log [<slot-id> | mm] tftp <ip-address> <filename> [oobm]`

Copies both the active and standby management modules' crash logs to a user-specified file. If no parameter is specified, files from all modules (management and interface) are concatenated.

slot-id: *retrieves the crash log from the module in the specified slot.*

mm: *retrieves the crash logs from both management modules and concatenates them.*

oobm: *For switches that have a separate out-of-band management port, specifies that the transfer will be through the out-of-band management interface. (Default is transfer through the data interface.)*

Enabling or Disabling the USB Port

This feature allows configuration of the USB port with either the CLI or SNMP.

To enable/disable the USB port with the CLI:

Syntax: usb-port
no usb-port

*Enables the USB port. The **no** form of the command disables the USB port and any access to the device.*

To display the status of the USB port:

Syntax: show usb-port

Displays the status of the USB port. It can be enabled, disabled, or not present.

```
ProCurve(config)# show usb-port
USB port status: enabled
USB port power status: power on      (USB device detected in port)
USB port reseal status: USB reseal not required
```

Figure A-17. Example of show usb-port Command Output on version K.13.59 and later

```
ProCurve(config)# show usb-port
USB port status: enabled
USB port power status: power on      (USB device detected in port)
```

Figure A-18. Example of show usb-port Command Output on version K.14.XX

One of the following messages indicates the presence or absence of the USB device:

- Not able to sense device in USB port
- USB device detected in port
- no USB device detected in port

The reseal status messages can be one of the following (K.13.XX only):

- undetermined USB reseal requirement
- USB reseal not required
- USB device reseal required for USB autorun

The autorun feature only works when a USB device is inserted and the USB port is enabled.

Behavior of Autorun When USB Port is Disabled

Software Versions K.13.XX Operation

When using software version K.13.58, if the USB port is disabled (no `usb-port` command), the USB autorun function does not work in the USB port until the USB port is enabled, the config file is saved, and the switch is rebooted. The 5 volt power to the USB port remains on even after the USB port has been disabled.

For software versions after K.13.58, the 5 volt power applied to the USB port is synchronized with the enabling of the USB port, that is, when the USB port is enabled, the 5 volts are supplied; when the USB port is disabled, the 5 volts are not supplied. For previous software versions the power was supplied continuously. The autorun function does not require a switch reboot, but the USB device must be inserted at least once after the port is enabled so that the switch recognizes that the device is present. If the USB device is inserted and then the USB port is enabled, the switch does not recognize that a USB device is present.

Software Version K.14.XX Operation

For software versions K.14.XX, the USB port can be disabled and enabled without affecting the autorun feature. When the USB port is enabled, the autorun feature activates if a USB device is already inserted in the USB port.

Power is synchronized with the enabling and disabling of USB ports as described above for K.13.59 and later software.

Using USB Autorun

USB autorun helps ease the configuration of ProCurve switches by providing a way to auto-execute CLI commands from a USB flash drive. Using this solution, you can create a command file (also known as an AutoRun file), write it to a USB storage device, and then execute the file simply by inserting the USB device in to the switch's 'Auxiliary Port'. The AutoRun file gets executed automatically when autorun is enabled on the switch, and can be designed for various purposes: for example, to configure the switch, to update software, or to retrieve diagnostic logs for troubleshooting purposes.

The overall USB autorun solution requires the following components:

- A ProCurve switch which can securely use USB autorun to load authorized configurations and write reporting information. This requires software versions K.13.01, T.13.01 or greater.
- The network management application *ProCurve Manager Plus* (PCM+). PCM+ is required to create a valid AutoRun file and view the results after the file has been executed on the switch.
- A non-proprietary USB flash drive.

Note

The ability to create a valid AutoRun file will be incorporated into an upcoming ProCurve Manager update. Refer to the ProCurve Manager documentation for details. For guidelines on using the USB port for basic file copy capabilities, see "Using USB to Transfer Files to and from the Switch" on page A-22.

How It Works

The general process for using USB Autorun is as follows (*steps 1, 2, and 7 require an upcoming update to PCM+ as described above*):

1. Create an AutoRun file using PCM+. Refer to the ProCurve Manager documentation for details.

Note

Creating the AutoRun file in PCM+, includes the following steps:

- a. specify the target device or devices.
- b. create the CLI script to be executed on the target device(s).
- c. determine if the file will be signed and/or encrypted.

- d. determine if the file will be ‘run once’ (moved to a ‘processed’ directory on execution) or ‘run many’ (kept in the root directory of the flash drive from where it can be executed again).

2. Deploy the AutoRun file to a USB flash drive.
3. (If required) Enable the autorun feature on the switch (autorun is enabled by default unless an operator or manager password has been set—see “Autorun and Configuring Passwords” on page A-51).
4. (If the AutoRun file has been signed or encrypted) Enable secure-mode on the switch firstly by configuring an encryption key and a valid trusted certificate, and then by enabling secure-mode via the CLI. See “Enabling Secure Mode” on page A-50.

5. Insert the USB flash drive into the switch’s USB auxiliary port.

The switch processes the AutoRun file automatically and writes a result (.txt) file and report (.xml) file back to the USB flash drive, reporting on the command operations that were executed.

6. Remove the USB device from the USB port.

The switch executes any post-commands, such as rebooting the switch to apply any configuration updates.

7. (Optional) Transfer the ‘result file’ and ‘report file’ to a PCM+-enabled computer for report checking. See “Troubleshooting Autorun Operations” on page A-49.

Security Considerations

By default, the switch is unsecured when shipped (that is, USB autorun is enabled by default). However, as soon as an operator or manager password is configured, autorun is disabled and must be re-enabled at the configuration level of the CLI before it can be used. The requirement to use PCM+ to create a valid AutoRun file helps prevent a non-authorized command file from being created and processed by the switch.

In terms of physical security, access to the switch’s console port and USB port are equivalent. Keeping the switch in a locked wiring closet or other secure space helps to prevent unauthorized physical access. As additional precautions, you have the following configuration options via the CLI (see page A-50):

- Disable autorun by setting an operator or manager password.
- Disable or re-enable the USB autorun function via the CLI.
- Enable autorun in secure mode to verify signatures in autorun command files and to decrypt encrypted command files.

Troubleshooting Autorun Operations

You can verify autorun operations by checking the following items:

USB Auxiliary Port LEDs. The following table shows LED indications on the Auxiliary Port that allow you to identify the different USB operation states.

Color	State	Meaning
Green	Slow Blinking	Switch is processing USB AutoRun file.
Green	Solid	Switch has finished processing USB AutoRun file. This LED state indicates the AutoRun file was successfully executed, and the report files were generated. The report files may be reviewed on a USB-enabled computer for more details. Upon removal of the USB device, the LED will be turned OFF.
n/a	Off	Indicates that no USB device has been inserted, or that a USB device that cannot be recognized as a USB storage device has been inserted, or that no AutoRun file can be found on the inserted USB device. If the USB device has just been removed from the port, the switch will execute any post commands.
Amber	Fast Blinking	Processing Error. The AutoRun file will stop processing when an error is encountered (for example, no more disk space is available on the USB device to write the result and report files). Remove the USB device and inspect its contents on a USB-enabled computer for more information on the error.

AutoRun Status Files. The following files are generated during autorun operations and written to the USB flash drive:

- Report file(s) (.xml file)—shows which CLI commands have been run. The file name includes a serial number and datetime stamp to indicate when and on which device the AutoRun file was executed.
- Result file(s) (.txt file)—contains the CLI output for each command that was run on the switch, allowing you to verify whether a command was executed successfully or not.

Note

PCM+ provides a mechanism to read these status files and capture the results of the commands executed. It also allows you to verify the report files for their authenticity and reject files that have not been signed (refer to the ProCurve Manager documentation for details).

The status files will not include any records of post commands that may have been executed after the USB flash drive was removed from the switch.

Event Log or Syslog. For details on how to use the switch's event log or syslog for help in isolating autorun-related problems, see "Using the Event Log for Troubleshooting Switch Problems" on page C-27.

Configuring Autorun on the Switch

To enable/disable the autorun feature on the switch, the following commands can be executed from configuration mode in the CLI.

Syntax: [no] autorun [encryption-key <key-string> | secure-mode]

Enables/disables USB autorun on the switch.

*Use the **encryption-key** keyword to configure or remove an encryption-key (a base-64 encoded string). The encryption key is a pre-requisite for enabling autorun in secure-mode. Encryption is regarded only when the AutoRun file is also signed by an authentic source.*

*Use the **secure-mode** keyword to enable or disable secure mode for autorun.*

Default: Enabled (or Disabled if a password has been set).

Enabling Secure Mode

Autorun secure mode can be used to verify the authenticity of autorun command files. Secure-mode is configured using the **autorun secure-mode** command and can be enabled under the following conditions:

- an encryption-key has already been configured using the **autorun encryption key** command; and
- a trusted certificate for verifying autorun command files has been copied to the switch using the **copy <tftp | usb> autorun-cert-file** command.

There is an additional security option to install a valid key-pair for signing the result files that are generated during autorun operations. The key-pair can be generated on the switch using the **crypto key generate autorun [rsa]** command.

Note

The key-pair can also be installed from a tftp server or via the usb port using **copy <tftp | usb> autorun-key-file <ipaddr filename>** command. The filename must contain the private key and the matching public key in a X509 certificate structure. Both the private key and the X509 certificate must be in PEM format.

Operating Notes and Restrictions

- Autorun is enabled by default, until passwords are set on the device.
- Secure-mode and encryption-key are disabled by default.
- To enable secure mode both an encryption key and trusted certificate must be set.
- If secure-mode is enabled, the following conditions apply:
 - the encryption-key cannot be removed/un-configured;
 - the key-pair cannot be removed.
- If secure mode is disabled, the key-pair can be removed using the **crypto key zeorize autorun** command.
- When installing the autorun certificate file and/or the other key files, the files must be in PEM format.

Autorun and Configuring Passwords

When an operator or manager password is configured on a switch, autorun will be disabled automatically, and a message is displayed on the screen as shown in the following example:

```
ProCurve# password manager
New password for manager: *****
Please retype new password for manager: *****
Autorun is disabled as operator/manager is configured.
```

After passwords are set, autorun can be re-enabled as needed using the **autorun** command.

For more information on configuring passwords, refer to the chapter on “Username and Password Security” in the *Access Security Guide* for your switch.

Viewing Autorun Configuration Information

The **show autorun** command displays autorun configuration status information as shown in the following example.

```
ProCurve(config)# show autorun

Autorun configuration status

Enabled           : Yes
Secure-mode      : Disabled
Encryption-key   :
```

Monitoring and Analyzing Switch Operation

Contents

Overview	B-3
Status and Counters Data	B-4
Menu Access To Status and Counters	B-5
General System Information	B-6
Menu Access	B-6
CLI Access to System Information	B-7
Task Monitor—Collecting Processor Data	B-8
Switch Management Address Information	B-9
Menu Access	B-9
CLI Access	B-10
Module Information	B-11
Menu: Displaying Port Status	B-11
CLI Access	B-12
Port Status	B-13
Menu: Displaying Port Status	B-13
CLI Access	B-14
Web Access	B-14
Viewing Port and Trunk Group Statistics and Flow Control Status	B-14
Menu Access to Port and Trunk Statistics	B-16
CLI Access To Port and Trunk Group Statistics	B-17
Web Browser Access To View Port and Trunk Group Statistics	B-18
Viewing the Switch's MAC Address Tables	B-18
Menu Access to the MAC Address Views and Searches	B-18
CLI Access for MAC Address Views and Searches	B-21
Spanning Tree Protocol (MSTP) Information	B-22
CLI Access to MSTP Data	B-22
Internet Group Management Protocol (IGMP) Status	B-23

Web Browser Interface Status Information	B-26
Traffic Mirroring	B-27
Mirroring Terminology	B-29
Mirrored Traffic Destinations	B-32
Local Destinations	B-32
Remote Destinations	B-32
Monitored Traffic Sources	B-32
Criteria for Selecting Mirrored Traffic	B-33
Mirroring Session Limits	B-33
Mirroring Sessions	B-33
Mirroring Configuration	B-34
Remote Mirroring Endpoint and Intermediate Devices	B-35
Migration to Release K.12.xx	B-36
Migration to Release K.14.01 or Greater	B-36
Using the Menu or Web Interface To Configure Local Mirroring ..	B-38
Menu and Web Interface Limits	B-38
Configuration Steps	B-39
CLI: Configuring Local and Remote Mirroring	B-42
Local Mirroring Overview	B-43
Remote Mirroring Overview	B-45
1. Determine the Mirroring Session and Destination	B-48
2. Configure a Mirroring Destination on a Remote Switch	B-49
3. Configure a Mirroring Session on the Source Switch	B-51
4. Configure the Monitored Traffic in a Mirror Session	B-54
Traffic Selection Options	B-54
Mirroring-Source Restrictions	B-55
Selecting All Inbound/Outbound Traffic to Mirror	B-56
Port Interface with Traffic Direction as the Selection Criteria	B-56
Untagged Mirrored Packets	B-58
VLAN Interface with Traffic Direction as the Selection Criteria	B-60
Selecting Inbound Traffic Using an ACL (Deprecated)	B-61
Selecting Inbound/Outbound Traffic Using a MAC Address	B-62
Selecting Inbound Traffic Using Advanced Classifier-Based Mirroring	B-65
Classifier-Based Mirroring Configuration	B-66

Viewing a Classifier-Based Mirroring Configuration	B-72
Classifier-Based Mirroring Restrictions	B-72
Applying Multiple Mirroring Sessions to an Interface	B-74
Displaying a Mirroring Configuration	B-75
Displaying All Mirroring Sessions Configured on the Switch .	B-75
Displaying the Remote Endpoints Configured on the Switch .	B-77
Displaying the Mirroring Configuration for a Specific Session	B-78
Displaying Resource Usage for Mirroring Policies	B-83
Viewing the Mirroring Configurations in the Running Configuration File	B-85
Mirroring Configuration Examples	B-86
Example: Local Mirroring Using Traffic-Direction Criteria . . .	B-86
Example: Remote Mirroring Using a Classifier-Based Policy .	B-87
Example: Remote Mirroring Using Traffic-Direction Criteria .	B-89
Maximum Supported Frame Size	B-91
Enabling Jumbo Frames To Increase the Mirroring Path MTU	B-92
Effect of Downstream VLAN Tagging on Untagged, Mirrored Traffic	B-93
Operating Notes for Traffic Mirroring	B-94
Troubleshooting Traffic Mirroring	B-96

Overview

The switches covered in this guide have several built-in tools for monitoring, analyzing, and troubleshooting switch and network operation:

- **Status:** Includes options for displaying general switch information, management address data, port status, port and trunk group statistics, MAC addresses detected on each port or VLAN, and STP, IGMP, and VLAN data (page B-5).
- **Counters:** Display details of traffic volume on individual ports (page B-15).
- **Event Log:** Lists switch operating events (“Using the Event Log for Troubleshooting Switch Problems” on page C-27).
- **Alert Log:** Lists network occurrences detected by the switch—in the Status | Overview screen of the web browser interface (page 5-21).
- **Configurable trap receivers:** Uses SNMP to enable management stations on your network to receive SNMP traps from the switch. (Refer to “SNMPv1 and SNMPv2c Traps” on page 14-19.)
- **Port monitoring (mirroring):** Copy all traffic from the specified ports to a designated monitoring port (page B-28).

Note

Link test and ping test—analysis tools in troubleshooting situations—are described in Appendix C, “Troubleshooting”. Refer to “Diagnostic Tools” on page C-62.

Status and Counters Data

This section describes the status and counters screens available through the switch console interface and/or the web browser interface.

Note

You can access all console screens from the web browser interface via Telnet to the console. Telnet access to the switch is available in the Device View window under the **Configuration** tab.

Status or Counters Type	Interface	Purpose	Page
Menu Access to Status and Counters	Menu	Access menu interface for status and counter data.	B-6
General System Information	Menu, CLI	Lists switch-level operating information.	B-7
Management Address Information	Menu, CLI	Lists the MAC address, IP address, and IPX network number for each VLAN or, if no VLANs are configured, for the switch.	B-10
Module Information	Menu, CLI	Lists the module type and description for each slot in which a module is installed.	B-12
Port Status	Menu, CLI, Web	Displays the operational status of each port.	B-14
Port and Trunk Statistics and Flow Control Status	Menu, CLI, Web	Summarizes port activity and lists per-port flow control status.	B-15
VLAN Address Table	Menu, CLI	Lists the MAC addresses of nodes the switch has detected on specific VLANs, with the corresponding switch port.	B-19
Port Address Table	Menu, CLI	Lists the MAC addresses that the switch has learned from the selected port.	B-19
STP Information	Menu, CLI	Lists Spanning Tree Protocol data for the switch and for individual ports. If VLANs are configured, reports on a per-VLAN basis.	B-23
IGMP Status	Menu, CLI	Lists IGMP groups, reports, queries, and port on which querier is located.	B-24
VLAN Information	Menu, CLI	For each VLAN configured in the switch, lists 802.1Q VLAN ID and up/down status.	B-25
Port Status Overview and Port Counters	Web	Shows port utilization and counters, and the Alert Log.	B-27

Menu Access To Status and Counters

Beginning at the Main Menu, display the Status and Counters menu by selecting:

1. Status and Counters

```
===== CONSOLE - MANAGER MODE =====  
Status and Counters Menu  
  
1. General System Information  
2. Switch Management Address Information  
3. Module Information  
4. Port Status  
5. Port Counters  
6. Vlan Address Table  
7. Port Address Table  
8. Spanning Tree Information  
0. Return to Main Menu...  
  
Displays switch management information including software versions.  
To select menu item, press item number, or highlight item and press <Enter>.
```

Figure B-1. The Status and Counters Menu

Each of the above menu items accesses the read-only screens described on the following pages. Refer to the online help for a description of the entries displayed in these screens.

General System Information

Menu Access

From the console Main Menu, select:

1. Status and Counters

1. General System Information

```
===== CONSOLE - MANAGER MODE =====
                        Status and Counters - General System Information

System Contact      :
System Location     :

Firmware revision   : K.11.00           Base MAC Addr      : 0001e7-a09900
ROM Version         : K.11.Z4           Serial Number      : S2600017409

Up Time             : 2 hours            Memory - Total     : 24,588,136
CPU Util (%)        : 1                  Free               : 19,613,568

IP Mgmt - Pkts Rx   : 0                  Packet - Total     : 832
           Pkts Tx   : 0                  Buffers  Free      : 793
                                           Lowest   : 769
                                           Missed   : 0
                                           24,588,1 6

Actions->  Back      Help
Return to previous screen.
Use arrow keys to change action selection and <Enter> to execute action.
```

Figure B-2. Example of General Switch Information

This screen dynamically indicates how individual switch resources are being used. Refer to the online Help for details.

CLI Access to System Information

The **show system** command displays general system information about the switch.

Syntax: show system [chassislocate | information | power-supply | temperature | fans]

Displays global system information and operational parameters for the switch.

chassislocate

Displays the chassisLocator LED status. Possible values are ON, Off, or Blink. When the status is On or Blink, the number of minutes that the Locator LED will continue to be on or to blink is displayed.

information

Displays global system information and operational parameters for the switch.

power-supply

Shows chassis power supply and settings.

temperature

Shows system temperature and settings.

fans

Shows system fan status.

```
ProCurve(config)# show system chassislocate
Chassis Locator LED: ON 5 minutes 5 seconds
ProCurve(config)# show system chassislocate
Chassis Locator LED: BLINK 10 minutes 6 seconds
ProCurve(config)# show system chassislocate
Chassis Locator LED: OFF
```

Figure B-3. Example of Command Results for show system chassislocate Command

```
ProCurve(config)# show system fans

Fan Information
  Num | State | Failures
-----+-----+-----
Sys-1 | Fan OK | 0

0 / 1 Fans in Failure State
0 / 1 Fans have been in Failure State
```

Figure B-4. Example of System Fan Status

```
ProCurve(config)# show system

Status and Counters - General System Information

System Name       : ProCurve Switch
System Contact    :
System Location   :

MAC Age Time (sec) : 300

Time Zone         : 0
Daylight Time Rule : None

Software revision : T.13.XX           Base MAC Addr      : 001635-b57cc0
ROM Version       : K.12.12         Serial Number      : LP621KI005

Up Time          : 51 secs           Memory - Total     : 152,455,616
CPU Util (%)     : 3                 Memory - Free      : 110,527,264

IP Mgmt - Pkts Rx : 0                 Packet - Total     : 6750
          Pkts Tx : 0                 Buffers - Free    : 5086
                                          Buffers - Lowest  : 5086
                                          Buffers - Missed  : 0
```

Figure B-5. Example of Switch System Information

Task Monitor—Collecting Processor Data

The task monitor feature allows you to enable or disable the collection of processor utilization data. The **task-monitor cpu** command is equivalent to the existing debug mode command “**taskusage -d**”. (The **taskUsageShow** command is available as well.)

When the **task-monitor** command is enabled, the **show cpu** command summarizes the processor usage by protocol and system functions.

Syntax: [no] task-monitor cpu

Allows the collection of processor utilization data. Only manager logins can execute this command. The settings are not persistent, that is, there are no changes to the configuration.

Default: Disabled

```
ProCurve(config)# task-monitor cpu
ProCurve(config)# show cpu

2 percent busy, from 2865 sec ago
1 sec ave: 9 percent busy
5 sec ave: 9 percent busy
1 min ave: 1 percent busy

% CPU | Description
-----+-----
    99 | Idle
```

Figure B-6. Example of the task-monitor cpu Command and show cpu Output

Switch Management Address Information

Menu Access

From the Main Menu, select:

1 Status and Counters ...

2. Switch Management Address Information

```
----- CONSOLE - MANAGER MODE -----  
Status and Counters - Management Address Information  
  
Time Server Address : Disabled  
  
VLAN Name      MAC Address      IP Address  
-----  
DEFAULT VLAN  0001e7-a09900   10.28.227.101  
VLAN-22       0001e7-a09900   Disabled  
VLAN-33       0001e7-a09900   Disabled  
  
Actions->    Back      Help  
Return to previous screen.  
Use arrow keys to change action selection and <Enter> to execute action.
```

Figure B-7. Example of Management Address Information with VLANs Configured

This screen displays addresses that are important for management of the switch. If multiple VLANs are *not* configured, this screen displays a single IP address for the entire switch. Refer to the online Help for details.

Note

As shown in figure B-7, all VLANs on the switches use the same MAC address. (This includes both the statically configured VLANs and any dynamic VLANs existing on the switch as a result of GVRP operation.)

Also, the switches covered in this guide use a multiple forwarding database. When using multiple VLANs and connecting a switch to a device that uses a single forwarding database, such as a Switch 4000M, there are cabling and tagged port VLAN requirements. For more on this topic, refer to the section titled “Multiple VLAN Considerations” in the “Static Virtual LANs (VLANs)” chapter of the *Advanced Traffic Management Guide* for your switch.

CLI Access

Syntax: show management

Module Information

Use this feature to determine which slots have modules installed and which type(s) of modules are installed.

Menu: Displaying Port Status

From the Main Menu, select:

1. Status and Counters ...
3. Module Information

```
ProCurve16-Dec-2005 16:29:21
----- CONSOLE - MANAGER MODE -----
                Status and Counters - Module Information

Slot           Module Description           Serial Number
-----
[A] ProCurve J8702A XL 24 port Gig-T POE module SG111sz235
C   ProCurve J8702A XL 24 port Gig-T POE module SG111sz236
D   ProCurve J8702 XL 4 port 10G X2 module      SG111sz237

Actions->  Back  Help

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

Figure B-8. Example of Module Information in the Menu Interface

CLI Access

The CLI **show modules** command will display additional component information for the following:

- System Support Modules (SSM)—identification, including serial number
- Mini-GBICS—a list of installed mini-GBICs displaying the type, “J” number, and serial number (when available)

Syntax: show modules [details]

Displays information about the installed modules, including:

- *The slot in which the module is installed*
- *The module description*
- *The serial number*
- *The System Support Module description, serial number, and status (8200zl switches only)*

Additionally, the part number (J number) and serial number of the chassis is displayed.

```
ProCurve(config)# show modules

Status and Counters - Module Information

Chassis: 5406zl J8697A          Serial Number:  SG560TN124
Slot  Module Description          Serial Number
-----
A     ProCurve J8706A 24p SFP zl Module  AD722BX88F
B     ProCurve J8702A 24p Gig-T zl Module FE999CV77F
C     ProCurve J8707A 4p 10-Gbe zl Module FB345DC99D
```

Figure B-9. Example of the show modules Command Output

```
ProCurve(config)# show modules details

Status and Counters - Module Information

Chassis: 8212zl J8715A          Serial Number:  SG560TN124
Slot  Module Description          Serial Number  Status
-----
MM1   ProCurve J9092A Management Module 8200zl  AD722BX88F    Active
SSM   ProCurve J8784A System Support Module  AF988DC78G    Active
C     ProCurve J8750A 20p +4 Mini-GBIC Module  446S2BX007    Active
      GBIC 1: J4859B 1GB LX-LC             4720347DFED734
      GBIC 2: J4859B 1GB LX-LC             4720347DFED735
```

Figure B-10. An Example of the show modules details Command for the 8212zl Showing SSM and Mini-GBIC Information

Note

On ProCurve 3500y1 and 6200y1 series switches, the mini-GBIC information does not display as the ports are fixed and not part of any module.

Port Status

The web browser interface and the console interface show the same port status data.

Menu: Displaying Port Status

From the Main Menu, select:

- 1. Status and Counters ...**
- 4. Port Status**


```

-----
                        Status and Counters - Port Status
-----
Port      Type      Intrusion      Enabled  Status  Mode      Flow
Alert
-----
A1                No             Yes      Down    Down    off
A2                No             Yes      Down    Down    off
A3                No             Yes      Down    Down    off
A4                No             Yes      Down    Down    off
B1      10/100TX  No             Yes      Up      100FDx  off
B2      10/100TX  No             Yes      Down    10FDx  off
B3      10/100TX  No             Yes      Down    10FDx  off
B4      10/100TX  No             Yes      Down    10FDx  off
B5      10/100TX  No             Yes      Down    10FDx  off
B6      10/100TX  No             Yes      Down    10FDx  off
B7      10/100TX  No             Yes      Down    10FDx  off

Actions->  Back      Intrusion log  Help

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

```

Figure B-11. Example of Port Status on the Menu Interface

CLI Access

Syntax: show interfaces brief

Web Access

1. Click on the **Status** tab.
2. Click on **[Port Status]**.

Viewing Port and Trunk Group Statistics and Flow Control Status

Feature	Default	Menu	CLI	Web
viewing port and trunk statistics for all ports, and flow control status	n/a	page B-17	page B-18	page B-19
viewing a detailed summary for a particular port or trunk	n/a	page B-17	page B-18	page B-19
resetting counters	n/a	page B-17	page B-18	page B-19

These features enable you to determine the traffic patterns for each port since the last reboot or reset of the switch. You can display:

- A general report of traffic on all LAN ports and trunk groups in the switch, along with the per-port flow control status (On or Off).
- A detailed summary of traffic on a selected port or trunk group.

You can also reset the counters for a specific port.

The menu interface and the web browser interface provide a dynamic display of counters summarizing the traffic on each port. The CLI lets you see a static “snapshot” of port or trunk group statistics at a particular moment.

As mentioned above, rebooting or resetting the switch resets the counters to zero. You can also reset the counters to zero for the current session. This is useful for troubleshooting. Refer to the “Note On Reset”, below.

Note on Reset

The **Reset** action resets the counter display to zero for the current session, but does not affect the cumulative values in the actual hardware counters. (In compliance with the SNMP standard, the values in the hardware counters are not reset to zero unless you reboot the switch.) Thus, using the **Reset** action resets the displayed counters to zero for the current session only. Exiting from the console session and starting a new session restores the counter displays to the accumulated values in the hardware counters.

Menu Access to Port and Trunk Statistics

To access this screen from the Main Menu, select:

1. Status and Counters ...


4. Port Counters

```

=====  CONSOLE - MANAGER MODE  =====
                        Status and Counters - Port Counters
-----
Port      Total Bytes  Total Frames  Errors Rx  Drops Tx  Flow
-----
A1        195,072      323           0           0      off
A2        651,816      871           0           0      off
A3-Trk1   290,163        500           0           0      off
A4-Trk1   260,134        501           0           0      off
C1        859,363        5147          0           0      off
C2        674,574        1693          0           0      off
C3        26,554         246           0           0      off
C4        113,184        276           0           0      off
C5         0             0             0           0      off
-----
Actions->  Back      Show details  Reset      Help
-----
Return to previous screen.
Use up/down arrow keys to scroll to other entries, left/right arrow keys to
change action selection, and <Enter> to execute action.

```

Figure B-12. Example of Port Counters on the Menu Interface

To view details about the traffic on a particular port, use the  key to highlight that port number, then select **Show Details**. For example, selecting port A2 displays a screen similar to figure B-13, below.

```

=====  CONSOLE - MANAGER MODE  =====
                        Status and Counters - Port Counters - Port A2
-----
Link Status      : up

Bytes Rx         : 630,746           Bytes Tx         : 21,070
Unicast Rx       : 568              Unicast Tx       : 285
Bcast/Mcast Rx   : 18              Bcast/Mcast Tx   : 0

PCS Rx           : 0                Drops Tx         : 0
Alignment Rx     : 0                Collisions Tx    : 0
Runts Rx         : 0                Late Colln Tx   : 0
Giants Rx        : 0                Excessive Colln : 0
Total Rx Errors  : 0                Deferred Tx      : 0

Actions->  Back      Reset      Help
-----
Return to previous screen.
Use arrow keys to change action selection and <Enter> to execute action.

```

Figure B-13. Example of the Display for Show details on a Selected Port

This screen also includes the **Reset** action for the current session. (Refer to the “Note on Reset” on page B-16.)

CLI Access To Port and Trunk Group Statistics

To Display the Port Counter Summary Report.

Syntax: show interfaces

This command provides an overview of port activity for all ports on the switch.

To Display a Detailed Traffic Summary for Specific Ports. .

Syntax: show interfaces <port-list>

This command provides traffic details for the port(s) you specify

To Reset the Port Counters.

It is useful to be able to clear all counters and statistics without rebooting the switch when troubleshooting network issues. The **clear statistics global** command clears all counters and statistics for all interfaces except SNMP. You can also clear the counters and statistics for an individual port using the **clear statistics <port-list>** command.

Syntax: clear statistics <<port-list> | global >

When executed with the <port-list> option, clears the counters and statistics for an individual port. When executed with the global option, clears all counters and statistics for all interfaces except SNMP.

The **show interfaces [<port-list>]** command displays the totals accumulated since the last boot or the last **clear statistics** command was executed. The menu and web pages also display these totals.

SNMP displays the counter and statistics totals accumulated since the last reboot; it is not affected by the **clear statistics global** command or the **clear statistics <port-list>** command. An SNMP trap is sent whenever the statistics are cleared.

Note

The clearing of statistics cannot be uncleared.

Web Browser Access To View Port and Trunk Group Statistics

1. Click on the **Status** tab.
2. Click on **[Port Counters]**.
3. To refresh the counters for a specific port, click anywhere in the row for that port, then click on **[Refresh]**.

Note

To reset the port counters to zero, you must reboot the switch.

Viewing the Switch's MAC Address Tables

Feature	Default	Menu	CLI	Web
viewing MAC addresses on all ports on a specific VLAN	n/a	page B-19	page B-22	—
viewing MAC addresses on a specific port	n/a	page B-21	page B-22	—
searching for a MAC address	n/a	page B-21	page B-22	—

These features help you to view:

- The MAC addresses that the switch has learned from network devices attached to the switch
- The port on which each MAC address was learned

Menu Access to the MAC Address Views and Searches

Per-VLAN MAC-Address Viewing and Searching. This feature lets you determine which switch port on a selected VLAN is being used to communicate with a specific device on the network. The per-VLAN listing includes:

- The MAC addresses that the switch has learned from network devices attached to the switch
- The port on which each MAC address was learned

1. From the Main Menu, select:
1. Status and Counters
5. VLAN Address Table
2. The switch then prompts you to select a VLAN.

```
Select VLAN : DEFAULT VLAN
```

3. Use the Space bar to select the VLAN you want, then press **[Enter]**. The switch then displays the MAC address table for that VLAN:

```
----- CONSOLE - MANAGER MODE -----  
Status and Counters - Address Table  
  
MAC Address   Located on Port  
-----  
0030c1-7f49c0 A3  
0030c1-7fec40 A1  
0030c1-b29ac0 A3  
0060b0-17de5b A3  
0060b0-880a80 A2  
0060b0-df1a00 A3  
0060b0-df2a00 A3  
0060b0-e9a200 A3  
009027-e74f90 A3  
080009-21ae84 A3  
080009-62c411 A3  
080009-6563e2 A3  
  
Actions-> Back   Search   Next page   Prev page   Help  
  
Return to previous screen.  
Use up/down arrow keys to scroll to other entries, left/right arrow keys to  
change action selection, and <Enter> to execute action.
```

Figure B-14. Example of the Address Table

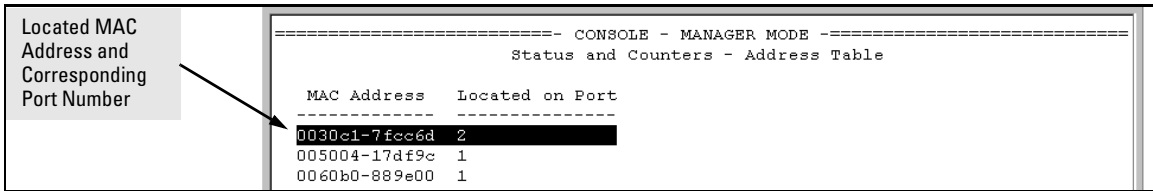
To page through the listing, use **Next page** and **Prev page**.

Finding the Port Connection for a Specific Device on a VLAN. This feature uses a device's MAC address that you enter to identify the port used by that device.

1. Proceeding from figure B-14, press **[S]** (for **Search**), to display the following prompt:

Enter MAC address: _

2. Type the MAC address you want to locate and press **[Enter]**. The address and port number are highlighted if found. If the switch does not find the MAC address on the currently selected VLAN, it leaves the MAC address listing empty.



```
===== CONSOLE - MANAGER MODE =====
                        Status and Counters - Address Table
-----
MAC Address      Located on Port
-----
0030e1-7fcc6d    2
005004-17df9c    1
0060b0-889e00    1
```

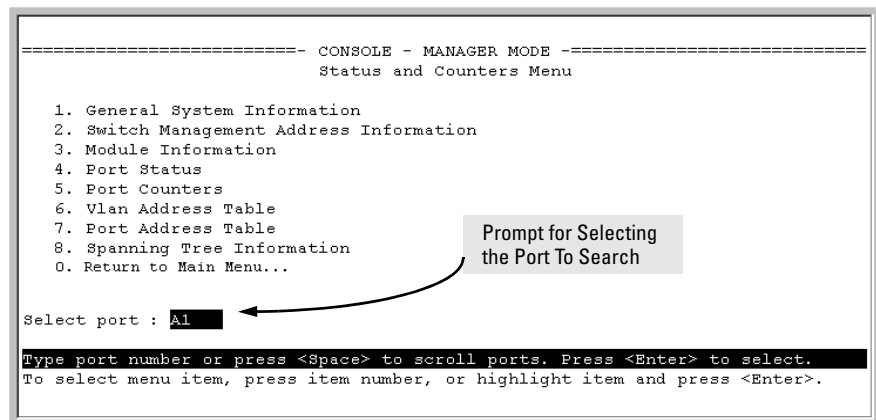
Figure B-15. Example of Menu Indicating Located MAC Address

3. Press **[P]** (for **Prev page**) to return to the full address table listing.

Port-Level MAC Address Viewing and Searching. This feature displays and searches for MAC addresses on the specified port instead of for all ports on the switch.

1. From the Main Menu, select:

- 1. Status and Counters**
 - 7. Port Address Table**



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

1. General System Information
2. Switch Management Address Information
3. Module Information
4. Port Status
5. Port Counters
6. Vlan Address Table
7. Port Address Table
8. Spanning Tree Information
0. Return to Main Menu...

Select port : A1
Type port number or press <Space> to scroll ports. Press <Enter> to select.
To select menu item, press item number, or highlight item and press <Enter>.
```

Figure B-16. Listing MAC Addresses for a Specific Port

2. Use the Space bar to select the port you want to list or search for MAC addresses, then press **[Enter]** to list the MAC addresses detected on that port.

Determining Whether a Specific Device Is Connected to the Selected Port. Proceeding from step 2, above:

1. Press **[S]** (for **S**earch), to display the following prompt:
Enter MAC address: _
2. Type the MAC address you want to locate and press **[Enter]**. The address is highlighted if found. If the switch does not find the address, it leaves the MAC address listing empty.
3. Press **[P]** (for **P**rev page) to return to the previous per-port listing.

CLI Access for MAC Address Views and Searches

Syntax: show mac-address
 [vlan <vlan-id >]
 [<port-list >]
 [<mac-addr >]

To List All Learned MAC Addresses on the Switch, with The Port Number on Which Each MAC Address Was Learned.

```
ProCurve> show mac-address
```

To List All Learned MAC Addresses on one or more ports, with Their Corresponding Port Numbers. For example, to list the learned MAC address on ports A1 through A4 and port A6:

```
ProCurve> show mac-address a1-a4,a6
```

To List All Learned MAC Addresses on a VLAN, with Their Port Numbers. This command lists the MAC addresses associated with the ports for a given VLAN. For example:

```
ProCurve> show mac-address vlan 100
```

Note

The switches covered in this guide operate with a multiple forwarding database architecture.

To Find the Port On Which the Switch Learned a Specific MAC Address. For example, to find the port on which the switch learns a MAC address of 080009-21ae84:


```
ProCurve# show mac-address 080009-21ae84
Status and Counters - Address Table - 080009-21ae84
  MAC Address : 080009-21ae84
  Located on Port : A2
```

Spanning Tree Protocol (MSTP) Information

CLI Access to MSTP Data

This option lists the MSTP configuration, root data, and per-port data (cost, priority, state, and designated bridge).

Syntax: show spanning-tree

This command displays the switch's global and regional spanning-tree status, plus the per-port spanning-tree operation at the regional level. Note that values for the following parameters appear only for ports connected to active devices: Designated Bridge, Hello Time, PtP, and Edge.

```
Switch-1(config)# show spanning-tree
Multiple Spanning Tree (MST) Information

STP Enabled      : Yes
Force Version    : MSTP-operation
IST Mapped VLANs : 1,66

Switch MAC Address : 0004ea-5e2000
Switch Priority    : 32768
Max Age          : 20
Max Hops         : 20
Forward Delay    : 15

Topology Change Count : 0
Time Since Last Change : 2 hours

CST Root MAC Address : 00022d-47367f
CST Root Priority     : 0
CST Root Path Cost   : 4000000
CST Root Port        : A1

IST Regional Root MAC Address : 000883-028300
IST Regional Root Priority     : 32768
IST Regional Root Path Cost   : 200000
IST Remaining Hops            : 19
```

Port	Type	Cost	Priority	State	Designated Bridge	Hello Time	PtP	Edge
A1	10/100TX	Auto	128	Forwarding	000883-028300	9	Yes	No
A2	10/100TX	Auto	128	Blocking	0001e7-948300	9	Yes	No
A3	10/100TX	Auto	128	Forwarding	000883-02a700	2	Yes	No
A4	10/100TX	Auto	128	Disabled				
A5	10/100TX	Auto	128	Disabled				
.				
.				

Figure B-17. Output from show spanning-tree Command

Internet Group Management Protocol (IGMP) Status

The switch uses the CLI to display the following IGMP status on a per-VLAN basis:

Show Command	Output
show ip igmp	Global command listing IGMP status for all VLANs configured in the switch: <ul style="list-style-type: none">• VLAN ID (VID) and name• Active group addresses per VLAN• Number of report and query packets per group• Querier access port per VLAN
show ip igmp <vlan-id>	Per-VLAN command listing above IGMP status for specified VLAN (VID)
show ip igmp group <ip-addr>	Lists the ports currently participating in the specified group, with port type, Access type, Age Timer data and Leave Timer data.

For example, suppose that **show ip igmp** listed an IGMP group address of 224.0.1.22. You could get additional data on that group by executing the following:

```
ProCurve> show ip igmp group 224.0.1.22

IGMP ports for group 224.0.1.22

  Port Type      Access      Age Timer  Leave Timer
  ----
  3    10/100TX  host        0          0
```

Figure B-18. Example of IGMP Group Data

VLAN Information

The switch uses the CLI to display the following VLAN status:

Show Command	Output
show vlan	Lists: <ul style="list-style-type: none"> • Maximum number of VLANs to support • Existing VLANs • Status (static or dynamic) • Primary VLAN
show vlan <vlan-id>	For the specified VLAN, lists: <ul style="list-style-type: none"> • Name, VID, and status (static/dynamic) • Per-Port mode (tagged, untagged, forbid, no/auto) • "Unknown VLAN" setting (Learn, Block, Disable) • Port status (up/down)

For example, suppose that your switch has the following VLANs:

Ports	VLAN	VID
A1-A12	DEFAULT_VLAN	1
A1, A2	VLAN-33	33
A3, A4	VLAN-44	44

The next three figures show how you could list data on the above VLANs.

Listing the VLAN ID (VID) and Status for ALL VLANs in the Switch.

```

ProCurve> show vlan
Status and Counters - VLAN Information
VLAN support : Yes
Maximum VLANs to support : 9
Primary VLAN: DEFAULT_VLAN

802.1Q VLAN ID Name          Status  __
-----
1          DEFAULT_VLAN  Static
33         VLAN-33      Static
44         VLAN-44      Static
```

Figure B-19. Example of VLAN Listing for the Entire Switch

Because ports A1 and A2 are not members of VLAN-44, it does not appear in this listing.

Listing the VLAN ID (VID) and Status for Specific Ports.

```
ProCurve>show vlan ports A1-A2
Status and Counters - VLAN Information - for ports A1,A2
802.1Q VLAN ID Name          Status
-----
1          DEFAULT_VLAN  Static
33         VLAN-33     Static
```

Figure B-20. Example of VLAN Listing for Specific Ports

Listing Individual VLAN Status.

```
ProCurve>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             Tagged   Learn      Up
A3             Untagged Learn      Up
A4             Untagged Learn      Down
A5             Untagged Learn      Down
*             *             *             *
*             *             *             *
*             *             *             *
```

Figure B-21. Example of Port Listing for an Individual VLAN

Web Browser Interface Status Information

The “home” screen for the web browser interface is the Status Overview screen, as shown below. As the title implies, it provides an overview of the status of the switch, including summary graphs indicating the network utilization on each of the switch ports, symbolic port status indicators, and the Alert Log, which informs you of any problems that may have occurred on the switch.

For more information on this screen, refer to the chapter titled “Using the ProCurve Web Browser Interface”.

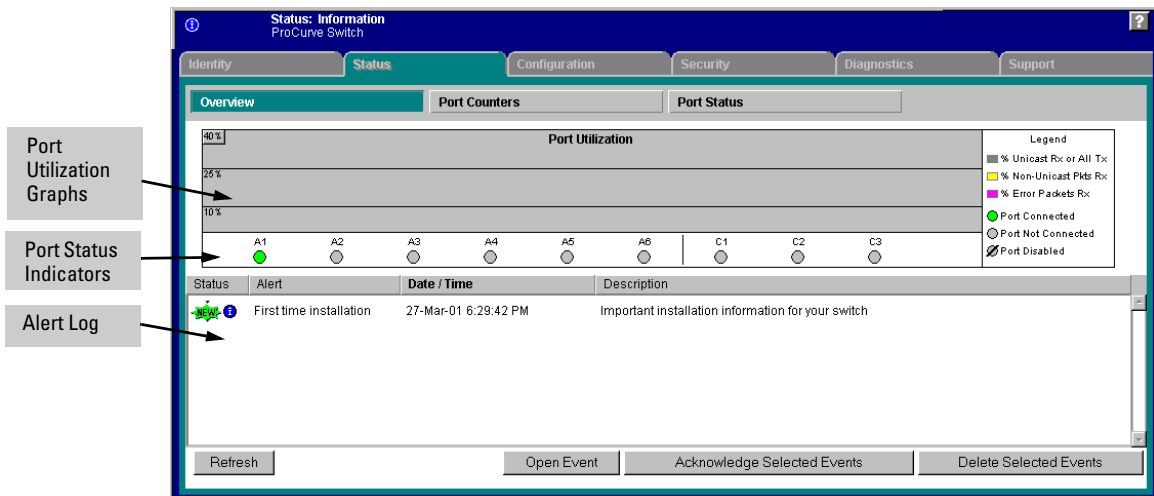


Figure B-22. Example of a Web Browser Interface Status Overview Screen

Traffic Mirroring

Mirror Features

Feature	Default	Menu	CLI
Mirror CLI Quick Reference	n/a	n/a	B-45, B-47
Configure Mirror Source	disabled	page B-39	page B-49
Configure Mirror Destination at Source	disabled	page B-39	page B-52
Configure Remote Mirroring at Destination	disabled	n/a	page B-50
Display Mirror Configuration	n/a	page B-39	page B-76

Starting in software release K.12.*xx*, traffic mirroring (Intelligent Mirroring) allows you to mirror (send a copy of) network traffic received or transmitted on a switch interface to a local or remote destination, such as a traffic analyzer or intrusion detection system (IDS).

Traffic mirroring provides the following benefits:

- Allows you to monitor the traffic flow on specific source interfaces.
- Helps in analyzing and debugging problems in network operation resulting from a misbehaving network or an individual client. The mirroring of selected traffic to an external device makes it easier to diagnose a network problem from a centralized location in a topology spread across a campus.
- Supports remote mirroring to simultaneously mirror switch traffic on one or more interfaces to multiple remote destinations. (In remote mirroring, you must first configure the remote mirroring endpoint — remote switch and exit port — before you specify a mirroring source for a session.)

Mirroring destinations. Traffic mirroring supports destination devices that are connected to the local switch or a remote switch:

- Traffic can be copied to a destination (host) device connected to the same switch as the mirroring source in a *local* mirroring session. You can configure up to *four* exit ports to which destination devices are connected in local mirroring sessions on a switch.
- Traffic can be bridged or routed to a destination device connected to a different switch in a *remote* mirroring session. You can configure up to 32 remote mirroring endpoints (IP address and exit port) to which destination devices are connected in remote mirroring sessions on a switch.

Mirroring sources and sessions. Traffic mirroring supports the configuration of port and VLAN interfaces as mirroring sources in up to *four* mirroring sessions on a switch. Each session can have one or more sources (ports and/or static trunks, a mesh, or a VLAN interface) that monitor traffic entering and/or leaving the switch.

Configuration Notes

Using the CLI, you can make full use of the switch's local and remote mirroring capabilities. Using the Menu interface, you can configure only local mirroring for either a single VLAN or a group of ports and/or static trunks.

Mirrored frames exceeding the allowed maximum transmission unit (MTU) size will be dropped. Also, the switch applies a 54-byte IPv4 header to mirrored frames. For more information, including the size limitation for jumbo and non-jumbo frames, see “Maximum Supported Frame Size” on page B-92.

Selecting mirrored traffic. You can use any of the following options to select the traffic to be mirrored on a port, trunk, mesh, or VLAN interface in a local or remote session:

- All traffic: Monitors all traffic entering or leaving the switch on one or more interfaces (*inbound and outbound*).
- Direction-based traffic selection: Monitors traffic that is either entering or leaving the switch (*inbound or outbound*). Monitoring traffic in only one direction improves operation by reducing the amount of traffic sent to a mirroring destination.
- MAC-based traffic selection: Monitors only traffic with a matching source and/or destination MAC address in packet headers entering and/or leaving the switch on one or more interfaces (*inbound and/or outbound*).
- Classifier-based service policy: Provides a finer granularity of match criteria to zoom in on a subset of monitored port or VLAN traffic (IPv4 or IPv6) and select it for local or remote mirroring (*inbound only*).

Deprecation of ACL-based Traffic Selection

In software release **K.14.01 or greater**, the use of ACLs for selecting traffic in a mirroring session has been deprecated and is replaced by the use of advanced classifier-based service policies (see “Selecting Inbound Traffic Using Advanced Classifier-Based Mirroring” on page B-66).

As with ACL criteria, classifier-based match/ignore criteria allow you to limit a mirroring session to selected inbound packets on a given port or VLAN interface (instead of mirroring all inbound traffic on the interface).

The following commands have been deprecated:

- **interface <port/trunk/mesh > monitor ip access-group <acl-name> in mirror < 1 - 4 | name-str >**
- **vlan < vid-# > monitor ip access-group <acl-name> in mirror < 1 - 4 | name-str >**

After you install and boot release K.14.01 or greater, ACL-based local and remote mirroring sessions configured on a port or VLAN interface are automatically converted to classifier-based mirroring policies. For more information, see “Migration to Release K.14.01 or Greater” on page B-37.

If you are running software release **K.13.x.x** or earlier, ACL permit/deny criteria are supported to select IP traffic entering a switch to mirror in a local or remote session, using specified source and/or destination criteria.

Mirroring Terminology

Figure B-23 shows an example of the terms used to describe the configuration of a sample local and remote mirroring session:

- In the local session, inbound traffic entering Switch A is monitored on port A2 and mirrored to a destination (host), traffic analyzer 1, through exit port A15 on the switch.

A local mirroring session means that the monitored interface (A2) and exit port (A15) are on the same switch.

- In the remote session, inbound traffic entering Switch A is monitored on port A1. A mirrored copy of monitored traffic is routed through the network to a remote mirroring endpoint: exit port B7 on Switch B. A destination device, traffic analyzer 2, is connected to the remote exit port.

A remote mirroring session means that:

- The monitored interface (A1) and exit port (B7) are on different switches.
- Mirrored traffic can be bridged or routed from a source switch to a remote switch.

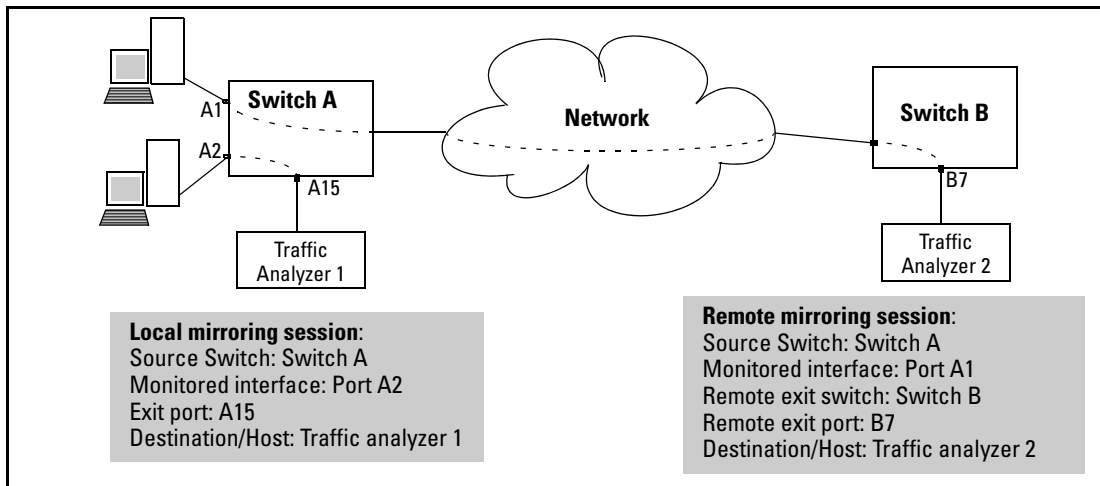


Figure B-23. Local and Remote Sessions Showing Mirroring Terms

Classifier-Based Mirroring Policy: The service policy applied to a monitored (port or VLAN) interface that specifies the classes of traffic to be copied to preconfigured mirroring destinations.

Destination : The host device that is connected to an exit port on the local source switch or a remote switch, and associated with a mirror-session number (1 to 4). See also *Exit Port* and *Host*.

Direction-Based Mirroring: On an interface configured for mirroring, the traffic direction (entering or leaving the switch, or both) is used as criteria for selecting the traffic to be mirrored.

Exit Port: The port to which a traffic analyzer or IDS is connected to receive mirrored traffic:

- For local mirroring, an exit port can be any port to which a traffic analyzer or IDS is connected and that is not configured as a monitored interface. You can configure up to four exit ports for local mirroring on a switch, using the command: **mirror <session> port <exit-port>**.
- For remote mirroring, the destination IP address (*dst-ip*) and exit port in a remote mirroring endpoint can belong to different VLANs. You can configure up to 32 exit ports for remote mirroring on a switch, using the command: **mirror endpoint ip <src-ip> <src-udp-port> <dst-ip> <exit-port>**.

Caution

An exit port should be connected only to a network analyzer, IDS, or other network edge device that has no connection to other network resources. Connecting a mirroring exit port to a network can result in serious network performance problems, and is strongly discouraged by ProCurve Networking.

Exit Switch: The switch with the exit port to which a destination device is connected. Depending on how mirroring is configured, the exit switch can be the local source switch or a remote switch. See also *Exit Port*.

Host: Used in this chapter to refer to a traffic analyzer or intrusion detection system (IDS).

IDS: Intrusion Detection System.

Local Mirroring: The monitored (source) interface and exit port in a mirroring session are on the same switch.

Monitored Interface: The interface (port, VLAN, trunk, or mesh) on the source switch on which the inbound and/or outbound traffic to be mirrored originates, configured with one of the **interface monitor** or **vlan monitor** commands (see “4. Configure the Monitored Traffic in a Mirror Session” on page B-55).

Remote Mirroring: The monitored (source) interface and exit port in a mirroring session are on different switches. For remote mirroring, you must always configure the IP destination address and exit port (the remote mirroring endpoint) before you configure the monitored interface, by using the following commands:

- On the remote (destination) switch:

mirror endpoint ip <src-ip> <src-udp-port> <dst-ip> <exit-port>

- On the local (source) switch:

mirror <session> **remote ip** <src-ip> <src-udp-port> <dst-ip>

For more information see *Exit Port* and “3. Configure a Mirroring Session on the Source Switch” on page B-52.

Source Switch: The source switch on which the inbound and/or outbound traffic to be mirrored originates. See also *Monitored Interface*.

Mirrored Traffic Destinations

Local Destinations

A local mirroring traffic destination is a port on the same switch as the source of the traffic being mirrored.

Remote Destinations

A *remote* mirroring traffic destination is a ProCurve switch configured to operate as the exit switch for mirrored traffic sessions originating on other ProCurve switches. As of June, 2007, switches capable of this operation include the following ProCurve switches:

- 3500yl
- 5400zl
- 6200yl
- 8200zl

Caution

After you configure a mirroring session with traffic-selection criteria and a destination, the switch immediately starts to mirror traffic to each destination device connected to an exit port. In a remote mirroring session which uses IPv4 encapsulation, if the intended exit switch is not already configured as the destination for the session, its performance may be adversely affected by the stream of mirrored traffic. For this reason, ProCurve strongly recommends that you configure the exit switch for a remote mirroring session before configuring the source switch for the same session.

Monitored Traffic Sources

You can configure mirroring for traffic entering or leaving the switch on:

- **Ports and static trunks:** Provides the flexibility for mirroring on individual ports, groups of ports, and/or static port trunks.
- **Meshed ports:** Enables traffic mirroring on all ports configured for meshing on the switch.
- **Static VLANs:** Supports traffic mirroring on static VLANs configured on the switch. This option enables easy mirroring of traffic from all ports on a VLAN. It automatically adjusts mirroring to include traffic from newly added ports, and to exclude traffic from ports removed from the VLAN.

Criteria for Selecting Mirrored Traffic

On the monitored sources listed above, you can configure the following criteria to select the traffic you want to mirror:

- Direction of traffic movement (entering or leaving the switch, or both)
- Type of IPv4 or IPv6 traffic entering the switch, as defined by a classifier-based service policy (see “Selecting Inbound Traffic Using Advanced Classifier-Based Mirroring” on page B-66)

In software release **K.14.01 or greater**, classifier-based service policies replace ACL-based traffic selection in mirroring sessions.

- Source and/or destination MAC addresses in packet headers

Mirroring Session Limits

A switch running software release K.12.*xxx* or greater supports the following:

- A maximum of four mirroring (local and remote) sessions
- A maximum of 32 remote mirroring endpoints (exit ports connected to a destination device that receive mirrored traffic originating from monitored interfaces on a different switch)

Mirroring Sessions

A mirroring session consists of a mirroring source and destination (endpoint). A mirroring source can be a port or static-trunk list, mesh, or VLAN interface. For any session, the destination must be a single (exit) port. The exit port cannot be a trunk, VLAN, or mesh interface.

Multiple mirroring sessions can be mapped to the same exit port, which provides flexibility in distributing hosts, such as traffic analyzers or an IDS. In a remote mirroring endpoint, the IP address of the exit port and the remote destination switch can belong to different VLANs.

Mirroring sessions can have the same or a different destination. You can configure an exit port on the local (source) switch and/or on a remote switch as the destination in a mirroring session. When configuring a mirroring destination, take into account the following options:

- Mirrored traffic belonging to different sessions can be directed to the same destination or to different destinations.

- You can reduce the risk of oversubscribing a single exit port by:
 - Directing traffic from different session sources to multiple exit ports
 - Configuring an exit port with a higher bandwidth than the monitored source port
- You can segregate traffic by type, direction, or source.

Mirroring Configuration

Table B-1 shows the different types of mirroring that you can configure using the CLI, Menu, and SNMP interfaces.

Table B-1. Mirroring Configuration Options

Monitoring Interface and Configuration Level	Traffic Selection Criteria	Traffic Direction		
		CLI Config	Menu and Web I/F Config ¹	SNMP Config
VLAN	All traffic	Inbound only Outbound only Both directions	All traffic (inbound and outbound combined)	Inbound only Outbound only Both directions
	ACL (IP traffic) ²	See “Selecting Inbound Traffic Using Advanced Classifier-Based Mirroring” on page B-67.		
	Classifier-based policy (IPv4 or IPv6 traffic)	Inbound only	Not available	Not available
Port(s) Trunk(s) Mesh	All traffic	Inbound only Outbound only Both directions	All traffic (inbound and outbound combined)	Inbound only Outbound only Both directions
	ACL (IP traffic) ²	See “Selecting Inbound Traffic Using Advanced Classifier-Based Mirroring” on page B-67.		
	Classifier-based policy (IPv4 or IPv6 traffic)	Inbound only	Not available	Not available
Switch (global)	MAC source/destination address	Inbound only Outbound only Both directions	Not available	Inbound only Outbound only Both directions
¹ Configures only session 1, and only for local mirroring. ² In release K.14.01 and greater, the use of ACLs to select inbound traffic in a mirroring session (using the <code><interface vlan> monitor ip access-group in mirror</code> command) has been deprecated and is replaced with classifier-based mirroring policies.				

Configuration Notes

Using the CLI, you can configure all mirroring options on a switch.

Using the Menu or Web interface, you can configure only session 1 and only local mirroring in session 1 for traffic in both directions on specified interfaces. (If session 1 has been already configured in the CLI for local mirroring for inbound-only or outbound-only traffic and you use the Menu or Web interface to modify the session 1 configuration, session 1 is *automatically* reconfigured to monitor both inbound and outbound traffic on the assigned interfaces. If session 1 has been configured in the CLI with a classifier-based mirroring policy or as a remote mirroring session, an error message is displayed if you try to use the Menu or Web interface to configure the session.)

You can use the CLI can configure sessions 1 to 4 for local or remote mirroring in any combination, and override a Menu or Web interface-based configuration of session 1.

You can also use SNMP configure sessions 1 to 4 for local or remote mirroring in any combination, and override a Menu or Web interface-based configuration of session 1, *except* that SNMP cannot be used to configure a classifier-based mirroring policy.

Remote Mirroring Endpoint and Intermediate Devices

The remote mirroring endpoint that is used in a remote mirroring session must be a ProCurve switch that supports the mirroring functions described in this chapter. (A remote mirroring endpoint consists of the remote switch and exit port connected to a destination device.) Because remote mirroring on a ProCurve switch uses IPv4 to encapsulate mirrored traffic sent to a remote endpoint switch, the intermediate switches and routers in a layer 2/3 domain can be from any vendor if they support IPv4.

The following restrictions apply to remote endpoint switches and intermediate devices in a network configured for traffic mirroring:

- The exit port for a mirroring destination must be an individual port, and *not* a trunk, mesh or VLAN interface.
- A switch mirrors traffic on static trunks, but *not* on dynamic LACP trunks.
- A switch mirrors traffic at line rate. When mirroring multiple interfaces in networks with high traffic levels, it is possible to copy more traffic to a mirroring destination than the link supports. However, some mirrored traffic may not reach the destination. If you are mirroring a high traffic volume, you can reduce the risk of oversubscribing a single exit port by:

- Directing traffic from different session sources to multiple exit ports
- Configuring an exit port with a higher bandwidth than the monitored source port

Migration to Release K.12.xx

On a switch that is running a software release earlier than K.12.xx with one or more mirroring sessions configured, when you download and boot release K.12.xx, the existing mirroring configurations are managed as follows:

- A legacy mirroring configuration on a port or VLAN interface maps to session 1.
- Traffic-selection criteria for session 1 is set to **both**; both inbound and outbound traffic (traffic entering *and* leaving the switch) on the configured interface is selected for mirroring.
- In a legacy mirroring configuration, a local exit port is applied to session 1.

Booting from Software Versions Earlier than K.12.xx: If it is necessary to boot the switch from a legacy (pre-K.12.xx) software version after using version K.12.xx or greater to configure mirroring, remove mirroring from the configuration before booting with the earlier software.

Maximum Supported Frame Size: The IPv4 encapsulation of mirrored traffic adds a 54-byte header to each mirrored frame. If a resulting frame exceeds the MTU (Maximum Transmission Unit) allowed in the path from the mirroring source to the mirroring destination, the frame is dropped. For more information, refer to “Maximum Supported Frame Size” on page B-92.

No Frame Truncation: Mirroring does not truncate frames, and oversized mirroring frames will be dropped. Also, remote mirroring does not allow downstream devices in a mirroring path to fragment mirrored frames.

Migration to Release K.14.01 or Greater

Note

If a switch is running software release K.12.xx, you must first upgrade to release K.13.xx before migrating the switch to release K.14.01 or greater.

When you download and boot software release K.14.01 or greater on a switch that is running release K.13.xx and has one or more mirroring sessions configured, an ACL-based mirroring configuration on a port or VLAN interface is mapped to a class and policy configuration based on the ACL.

The new mirroring policy is automatically configured on the same port or VLAN interface on which the mirroring ACL was assigned. The behavior of the new class and mirroring-policy configuration exactly matches the traffic-selection criteria and mirroring destination used in the ACL-based session.

Figures B-24 and B-25 show how ACL-based selection criteria in a mirroring session are converted to a classifier-based policy and class configuration when you install release K.14.01 or greater on a switch.

```
ProCurve(config)# show run
Running configuration:
. . .
ip access-list extended "100"
  10 permit icmp 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255 0
  exit
. . .
mirror 1 port C1
mirror 2 name "test-10" remote ip 10.10.10.1 8010 10.10.30.2
. . .
interface C1
  monitor ip access-group "100" In mirror 1
  exit
. . .
```

Configuration of ACL 100 that is used to select mirrored traffic in session 1

Existing mirror sessions configured on the switch for a local (port C1 in session 1) and remote (session 2) monitored interface

ACL-based traffic selection on monitored interface C1 in session 1

Figure B-24. Mirroring Configuration in “show run” Output in Release K.13.xx

```
ProCurve(config)# show run
Running configuration:
. . .
mirror 1 port B3
mirror 2 name "test-10" remote ip 10.10.10.1 8010 10.10.30.2
. . .
class ipv4 "100MirrorClass"
  10 match icmp 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255 0
  exit
policy mirror "100MirrorPolicy"
  10 class ipv4 "100" action mirror 1
  exit
. . .
interface C1
  service-policy "100MirrorPolicy" In
  exit
. . .
```

After migration to release K.14.01 or greater, the existing mirroring configurations for sessions 1 (local) and 2 (remote) on the switch remain the same.

The traffic-selection criteria in ACL 100 (Figure B-24) applied to inbound traffic on port C1 in session 1 are converted to a class and policy configuration with the names, “100MirrorClass” and “100MirrorPolicy”, which are applied to inbound traffic on port C1 in session 1 with the **service-policy** command.

Figure B-25. Mirroring Configuration in “show run” Output in Release K.14.01 or Greater

Using the Menu or Web Interface To Configure Local Mirroring

Menu and Web Interface Limits

The Menu and Web interfaces can be used to quickly configure or reconfigure local mirroring on session 1, and allow one of the following two mirroring source options:

- any combination of source port(s), trunk(s), and/or a mesh
- one static, source VLAN interface

The Menu and Web interfaces also have these limits:

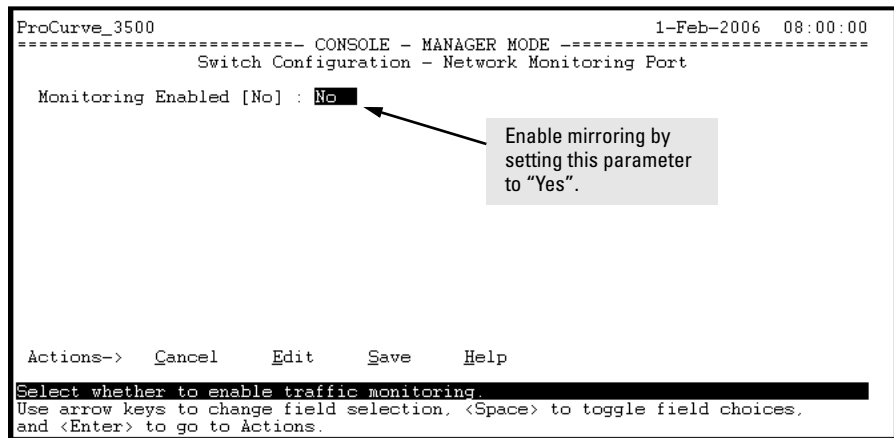
- Configure and display only session 1 and only as a local mirroring session for traffic in *both* directions on the specified interface. (Selecting inbound-only or outbound-only is not an option.)
- If session 1 has been configured in the CLI for local mirroring for inbound-only or outbound-only traffic on one or more interfaces, then using the Menu or Web interface to change the session 1 configuration *automatically reconfigures the session* to monitor both inbound and outbound traffic on the designated interface(s).
- If session 1 has been configured in the CLI with an ACL/classifier-based mirroring policy or as a remote mirroring session, then the Menu and Web interfaces are not available for changing the session 1 configuration.
- The CLI (and SNMP) can be used to override any Menu or Web interface configuration of session 1.

Configuration Steps

Notes

If mirroring has already been enabled on the switch, the Menu screens will appear differently than shown in this section.

1. From the Main Menu, Select:
 2. **Switch Configuration...**
 3. **Network Monitoring Port**



```
ProCurve_3500                                     1-Feb-2006 08:00:00
-----
----- CONSOLE - MANAGER MODE -----
Switch Configuration - Network Monitoring Port

Monitoring Enabled [No] : No

Actions->  Cancel      Edit      Save      Help

Select whether to enable traffic monitoring.
Use arrow keys to change field selection, <Space> to toggle field choices,
and <Enter> to go to Actions.
```

Figure B-26. The Default Network Mirroring Configuration Screen

2. In the Actions menu, press [E] (for Edit).
3. If mirroring is currently disabled for session 1 (the default), then enable it by pressing the Space bar (or [Y]) to select Yes.
4. Press the down arrow key to display a screen similar to the following and move the cursor to the **Monitoring Port** parameter.

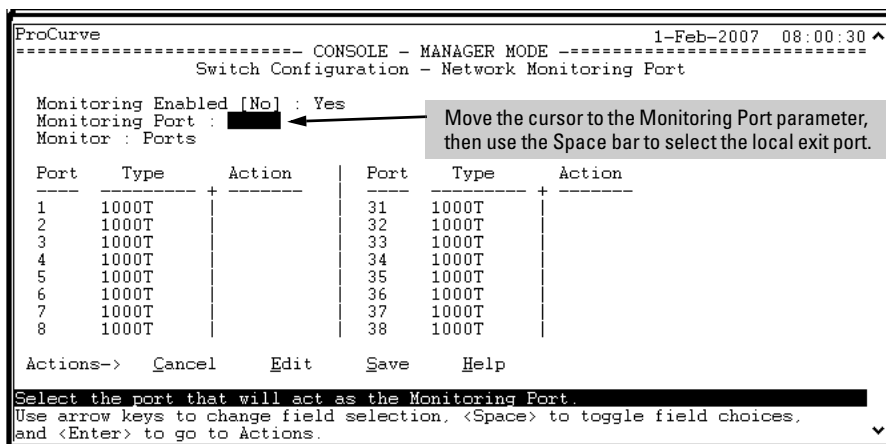



Figure B-27. How To Select a Local Exit Port

5. Use the Space bar to select the port to use for sending mirrored traffic to a locally connected traffic analyzer or IDS. (The selected interface must be a single port. It cannot be a trunk or mesh.) In this example, port 5 is selected as the local exit port.
6. Highlight the Monitor field and use the Space bar to select the interfaces to mirror:
 - Ports:** Use for mirroring ports, static trunks, or the mesh.
 - VLAN:** Use for mirroring a VLAN.
7. Do one of the following:
 - If you are mirroring ports, static trunks, or the mesh, go to step 8.
 - If you are mirroring a VLAN:
 - i. Press **[Tab]** or the down arrow key to move to the **VLAN** field.



- ii. Use the Space bar to select the VLAN you want to mirror.
- iii. Go to step 10.

- Use the down arrow key to move the cursor to the **Action** column for the individual port interfaces and position the cursor at a port, trunk, or mesh you want to mirror.

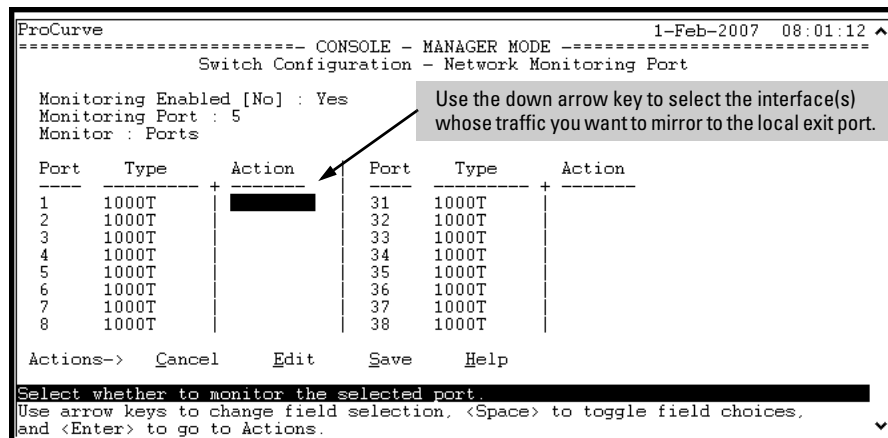
```
ProCurve 1-Feb-2007 08:01:12 ^
----- CONSOLE - MANAGER MODE -----
Switch Configuration - Network Monitoring Port

Monitoring Enabled [No] : Yes
Monitoring Port : 5
Monitor : Ports

Port  Type  Action  Port  Type  Action
-----+-----+-----+-----+-----+-----
1    1000T
2    1000T
3    1000T
4    1000T
5    1000T
6    1000T
7    1000T
8    1000T
31   1000T
32   1000T
33   1000T
34   1000T
35   1000T
36   1000T
37   1000T
38   1000T

Actions->  C_a_n_c_e_l  E_d_i_t  S_a_v_e  H_e_l_p

Select whether to monitor the selected port.
Use arrow keys to change field selection, <Space> to toggle field choices,
and <Enter> to go to Actions.
```



- Press the Space bar to select **Monitor** for the port(s) and/or trunk(s) and/or mesh that you want mirrored. Use the down arrow key to move from one interface to the next in the **Action** column. (If the mesh or any trunks are configured, they will appear at the end of the port listing.)
- When you finish selecting interfaces to mirror, press **[Enter]**, then press **[S]** (for **Save**) to save your changes and exit from the screen.
- Return to the Main Menu.

CLI: Configuring Local and Remote Mirroring

Command	Page
Quick Reference	
Local Mirroring Commands	B-45
Remote Mirroring Commands	B-47
Configuring a Remote Mirroring Destination	
On the remote switch: mirror endpoint ip < src-ip > < src-udp-port > < dst-ip > < exit-port >	B-50
On the local switch: mirror < session > remote ip < src-ip > < src-udp-port > < dst-ip >	B-53
Configuring a Local Mirroring Destination	
On the local switch: mirror < session > port < exit-port >	B-53
Configuring Monitored Traffic¹	
interface < port/trunk/mesh >	
monitor all < in out both > mirror < session > [no-tag-added]	B-57
monitor ip access-group < acl-name > in mirror < session > (deprecated)	B-62
service-policy < mirror-policy-name > in	B-67
vlan < vid-# >	B-61
monitor all < in out both > mirror < session >	B-61
monitor ip access-group < acl-name > in mirror < session > (deprecated)	B-63
service-policy < mirror-policy-name > in	B-67
monitor mac < mac-addr > < src dest both > mirror	
Display Commands	
show monitor [endpoint < session-number > name < session-name >]	B-76
Mirroring Examples	
	B-87
Maximum Frame Size	
	B-92
Operating Notes	
	B-95
¹ In release K.14.01 and greater, the use of ACLs to select inbound traffic in a mirroring session <interface vlan> monitor ip access-group in mirror command has been deprecated and is replaced with classifier-based mirroring policies.	

Using the CLI, you can configure a mirroring session for a destination device connected to an exit port on either:

- The same switch as the source interface (local mirroring)
- A different switch (remote mirroring). The remote switch must be a ProCurve switch offering the full mirroring capabilities described in this chapter.

For an overview of the procedures for configuring a local or remote mirroring session, refer to the following sections:

- “Local Mirroring Overview” on page B-44
- “Remote Mirroring Overview” on page B-46 (The remote switch must be a ProCurve switch offering the full mirroring capabilities described in this chapter.)

For a detailed description of each step in a mirroring configuration, refer to:

- “1. Determine the Mirroring Session and Destination” on page B-49
- “2. Configure a Mirroring Destination on a Remote Switch” on page B-50
- “3. Configure a Mirroring Session on the Source Switch” on page B-52
- “4. Configure the Monitored Traffic in a Mirror Session” on page B-55:
 - “Selecting All Inbound/Outbound Traffic to Mirror” on page B-57
 - “Selecting Inbound Traffic Using an ACL (Deprecated)” on page B-62 (Deprecated in release K.14.01 and greater)
 - “Selecting Inbound/Outbound Traffic Using a MAC Address” on page B-63
 - “Selecting Inbound Traffic Using Advanced Classifier-Based Mirroring” on page B-66

Caution

After you configure a mirroring session with traffic-selection criteria and a destination, the switch immediately starts to mirror traffic to each destination device connected to an exit port. In a remote mirroring session which uses IPv4 encapsulation, if the exit switch is not already configured as the destination for the session, its performance may be adversely affected by the stream of mirrored traffic. For this reason, ProCurve strongly recommends that you configure the exit switch for a remote mirroring session before configuring the source switch for the same session.

Local Mirroring Overview

To configure a local mirroring session in which the mirroring source and destination are on the same switch, follow these general steps:

1. Determine the session and local destination port:
 - Session number (1-4) and (optional) alphanumeric name
 - Exit port (any port on the switch except a monitored interface used to mirror traffic)
2. Enter the **mirror < session-# > [name < session-name >] port < port-# >** command to configure the session.

3. Determine the traffic to be selected for mirroring by any of the following methods and the appropriate configuration level (VLAN, port, mesh, trunk, switch):
 - a. Direction: inbound, outbound, or both
 - b. Classifier-based mirroring policy: inbound only for IPv4 or IPv6 traffic
 - c. MAC source and/or destination address: inbound, outbound, or both
4. Enter the **monitor** command to assign one or more source interfaces to the session.

After you complete step 4, the switch begins mirroring traffic to the configured exit port. The next two sections provide a quick reference to the configuration commands for local and remote mirroring sessions.

Quick Reference to Local Mirroring Set-Up. The following commands configure mirroring for a local session in which the mirroring source and destination are on the same switch. For command syntax details, refer to the pages listed with each heading.

- The **mirror** command identifies the destination in a mirroring session.
- The **interface** and **vlan** commands identify the mirroring source, including source interface, traffic direction, and traffic-selection criteria for a specified session.

Configure a Local Mirroring Session (Page B-52):

Mirror-Session Number, Local Exit Port, and (Optional) Session Name

```
[no] mirror < 1 - 4 > port < exit-port-# > [ name < name-str > ]
```

The **no mirror <session-#> port** command removes the mirroring session and any mirroring source previously assigned to that session by the following commands.

Configure Traffic-Direction Criteria to Select Traffic (Page B-57)

```
[no] < interface < port/trunk/mesh > | vlan < vid-# >>  
monitor all < in | out | both > mirror < session > [< session > ... ] [no-tag-added]
```

**Deprecated
Command**

Configure ACL Criteria to Select Inbound Traffic — Deprecated

```
[no] < interface < port/trunk/mesh > | vlan < vid-# >>  
monitor ip access-group < acl-name > in mirror < session > [< session > ... ]
```

Configure a Mirroring Policy to Select Inbound Traffic (Page B-66)

```
class < ipv4 | ipv6 > < classname >  
    [no] [seq-number] < match | ignore > < ip-protocol > < source-address >  
    < destination-address > [ precedence precedence-value ] [ tos tos-value ]  
    [ ip-dscp codepoint ] [ vlan vlan-id ]  
policy mirror < policy-name >  
    [no] [seq-number] class < ipv4 | ipv6 > <classname> action mirror <ses-  
sion> [action mirror < session > ... ]  
    [no] default-class action mirror < session-# >  
[no] < interface < port/trunk > | vlan < vid-# > service-policy <mirror-policy-name> in
```

In the **policy mirror** command, the **mirror < session >** parameter accepts a number (1 to 4) or name, if the specified mirroring session has already been configured with the **name < name-str >** option in the **mirror** command.

The **no < interface | vlan > service-policy in** command removes the mirroring policy from a port, VLAN, trunk, or mesh interface for a specified session, but leaves the session available for other assignments.

Configure MAC-based Criteria to Select Traffic (Page B-63)

```
[no] monitor mac < mac-addr > < src | dst | both > mirror < session >
```

Enter the **monitor mac mirror** command at the global configuration level.

Use the **no** form of the complete command syntax (for example, **no monitor mac 112233-445566 src mirror 3**) to remove a MAC address as mirroring criteria from an active session on the switch without removing the session itself.

Remote Mirroring Overview

To configure a remote mirroring session in which the mirroring source and destination are on different switches, follow these general steps:

1. Determine the IP addressing, UDP port number, and destination (exit) port number for the remote session:
 - a. Source VLAN or subnet IP address on the source switch
 - b. Destination VLAN or subnet IP address on the destination switch
 - c. Random UDP port number for the session (7933-65535)
 - d. Remote mirroring endpoint: Exit port and IP address of the remote destination switch (In a remote mirroring endpoint, the IP address of the exit port and remote switch can belong to different VLANs.)

Requirement: For remote mirroring, the same IP addressing and UDP port number must be configured on both the source and destination switches.

2. On the remote *destination* (endpoint) switch, enter the **mirror endpoint** command with the information from step 1 to configure a mirroring session for a specific exit port.
3. Determine the session (1 - 4) and (optional) alphanumeric name to use on the *source* switch.
4. Determine the traffic to be filtered by any of the following selection methods and the appropriate configuration level (VLAN, port, mesh, trunk, global):
 - a. Direction: inbound, outbound, or both
 - b. Classifier-based mirroring policy: inbound only for IPv4 or IPv6 traffic
 - c. MAC source and/or destination address: inbound, outbound, or both
5. On the *source* switch:
 - a. Enter the **mirror** command with the session number (1 - 4) and the IP addresses and UDP port number from step 1 to configure a mirroring session.
 - b. Enter one of the following commands to configure one or more of the traffic-selection methods in Step 4 for the configured session:
interface <port/trunk/mesh> <monitor | service-policy policy-name in >
vlan <vid> <monitor | service-policy policy-name in >
monitor mac <mac-addr >

After you complete Step 5b, the switch begins mirroring traffic to the remote destination (endpoint) configured for the session.

Quick Reference to Remote Mirroring Set-Up. The following commands configure mirroring for a remote session in which the mirroring source and destination are on different switches:

- The **mirror** command identifies the destination in a mirroring session.
- The **interface** and **vlan** commands identify the monitored interface, traffic direction, and traffic-selection criteria for a specified session.

Caution

When configuring a remote mirroring session, always configure the destination switch first. Configuring the source switch first can result in a large volume of mirrored, IPv4-encapsulated traffic arriving at the destination without an exit path, which can slow switch performance.

**Configure the Mirroring Destination on a Remote Switch (Page B-50):
IP Address and UDP Port on Source Switch
IP Address and Exit Port on Remote Switch**

```
mirror endpoint ip < src-ip-addr > < src-udp-port > < dst-ip-addr > port < exit-port >
```

Enter this command on a remote switch to configure the exit port to use in a remote mirroring session. You will configure the mirroring source on the local switch in the next step.

The **mirror endpoint ip** command configures:

- The unique UDP port number to be used for the mirroring session on the source switch. The recommended port range is from 7933 to 65535.
- The IP address of the source switch to use in the session
- The IP address and exit-port number on the remote (endpoint) switch

In a remote mirroring endpoint, the IP address of exit port and the remote destination switch can belong to different VLANs.

Configure the Mirroring Source on the Local Switch (Page B-52)

```
mirror < 1 - 4 > [name < name-str >] remote ip < src-ip > < src-udp-port > < dst-ip >
```

The **no mirror <1 - 4>** command form removes both the mirroring session and any mirroring source(s) previously assigned to the session by the following commands.

Configure Traffic-Direction Criteria to Select Traffic (Page B-57)

```
[no] < interface < port/trunk/mesh > | vlan < vid-# >>  
monitor all < in | out | both > mirror < 1 - 4 | name-str > [< 1 - 4 | name-str > ... ]
```

**Deprecated
Command**

Configure ACL Criteria to Select Inbound Traffic (Page B-62)

```
[no] < interface < port/trunk/mesh > | vlan < vid-# >>  
monitor ip access-group < acl-name > in  
mirror < 1 - 4 | name-str > [< 1 - 4 | name-str > ... ]
```

Configure a Mirroring Policy to Select Inbound Traffic (Page B-66)

```
class < ipv4 | ipv6 > < classname >
    [no] [seq-number] < match | ignore > < ip-protocol > < source-address >
    < destination-address > [ precedence precedence-value ] [ tos tos-value ]
    [ ip-dscp codepoint ] [ vlan vlan-id ]
policy mirror < policy-name >
    [no] [seq-number] class < ipv4 | ipv6 > < classname > action mirror <session>
    [ action mirror < session > ... ]
    [no] default-class action mirror < session >
[no] < interface < port/trunk > | vlan < vid-# > > service-policy <mirror-policy-name> in
```

In the **policy mirror** command, the **mirror < session >** parameter accepts a number (1 to 4) or name, if the specified mirroring session has already been configured with the **name < name-str >** option in the **mirror** command.

The **no < interface | vlan > service-policy in** command removes the mirroring configuration from a port, VLAN, trunk, or mesh interface for a specified session, but leaves the session available for other assignments.

Configure the MAC-based Criteria to Select Traffic (Page B-63)

```
[no] monitor mac < mac-addr > < src | dst | both > mirror < session >
```

Note

If you have already configured session 1 with a destination, you can enter the **vlan < vid > monitor** or **interface < port > monitor** command without traffic-selection criteria and session identifier to:

- Overwrite the existing session 1 configuration.
 - Automatically configure mirroring in session 1 for inbound and outbound traffic on specified VLAN or port interfaces with the preconfigured destination.
-

1. Determine the Mirroring Session and Destination

For a Local Mirroring Session. Determine the port number for the exit port (such as A5, B10, etc.), then go to “4. Configure the Monitored Traffic in a Mirror Session” on page B-55.

For a Remote Mirroring Session. Determine the following information and then go to step 2, below.

- The IP address of the VLAN or subnet on which the exit port exists on the destination switch
 - The port number of the remote exit port on the remote destination switch (In a remote mirroring endpoint, the IP address of the exit port and the remote destination switch can belong to different VLANs.)
-

- The IP address of the VLAN or subnet on which the mirrored traffic enters or leaves the source switch
- The unique UDP port number to use for the session on the source switch (The recommended port range is from 7933 to 65535.)

Caution

Although the switch supports the use of UDP port numbers from 1 to 65535, UDP port numbers below 7933 are reserved for various IP applications. Using these port numbers for mirroring can result in an interruption of other IP functions, and in non-mirrored traffic being received on the destination (end-point) switch and sent to the device connected to the remote exit port.

2. Configure a Mirroring Destination on a Remote Switch

This step is required only if you are configuring a remote mirroring session in which the exit port is on a different switch than the monitored (source) interface. If you are configuring local mirroring, go to step 3 on page B-55.

For remote mirroring, you must configure the *destination* switch to recognize each mirroring session and forward mirrored traffic to an exit port before you configure the *source* switch. Configure the destination switch with the values you determined for remote mirroring in “1. Determine the Mirroring Session and Destination” on page B-49.

Note

A remote destination switch can support up to 32 remote mirroring endpoints (exit ports connected to a destination device in a remote mirroring session).

Configuring a Destination Switch in a Remote Mirroring Session.

Enter the **mirror endpoint ip** command on the remote switch to configure the switch as a remote endpoint for a mirroring session with a different source switch.

Caution

When configuring a remote mirroring session, always configure the destination switch first. Configuring the source switch first can result in a large volume of mirrored, IPv4-encapsulated traffic arriving at the destination without an exit path, which can slow switch performance.

Syntax mirror endpoint ip < src-ip > < src-udp-port > < dst-ip > < exit-port-# >
no mirror endpoint ip < src-ip > < src-udp-port > < dst-ip >

This command is used on a destination switch to configure the remote endpoint of a mirroring session. The command uniquely associates the mirrored traffic from the desired session on a monitored source with a remote exit port on the destination switch. You must use the same set of source and destination parameters you when configure the same session on both the source and destination switches.

*For a given mirroring session, the same <src-ip >, <src-udp-port >, and <dst-ip > values must be entered with the **mirror endpoint ip** command on the destination switch, and later with the **mirror remote ip** command on the source switch. Refer to the **mirror remote ip** command syntax on page B-53 for more information.*

*The **no** form of the command deletes the mirroring endpoint for the configured session on the remote destination switch.*

Caution: *Do not remove the configuration of a remote mirroring endpoint support for a given session if there are source switches currently configured to mirror traffic to the endpoint.*

< src-ip >: *This parameter must exactly match the < src-ip > address you configure on the source switch for the remote session.*

Syntax: mirror endpoint ip < src-ip > < src-udp-port > < dst-ip > < exit-port-# >
no mirror endpoint ip < src-ip > < src-udp-port > < dst-ip >

— Continued —

< src-udp-port >: *This parameter must exactly match the < src-udp-port > value you configure on the source switch for the remote session. The recommended port range is 7933 to 65535.*

This setting associates the monitored source with the desired remote endpoint in the remote session by using the same, unique UDP port number to identify the session on the source and remote switches.

< dst-ip >: *This parameter must exactly match the < dst-ip > setting you configure on the source switch for the remote session.*

< exit-port-# >: *Exit port for mirrored traffic in the remote session, to which a traffic analyzer or IDS is connected.*

3. Configure a Mirroring Session on the Source Switch

To configure local mirroring, only a session number and exit port number are required. See “Configuring a Source Switch in a Local Mirroring Session” below for more information.

If the exit port for a mirroring destination is on a remote switch instead of the local (source) switch, then you must enter the source IP address, destination IP address, and UDP port number for the remote mirroring session (see page B-53).

Configuring a Source Switch in a Local Mirroring Session. For a local mirroring session, enter the **mirror port** command on the source switch to configure an exit port on the same switch. To create the mirroring session, use the information gathered in “1. Determine the Mirroring Session and Destination” on page B-49.

Syntax: mirror < 1 - 4 > port < exit-port-# > [name < name-str >]
no mirror < 1 - 4 >

This command assigns the exit port to use for the specified mirroring session, and must be executed from the global configuration level.

*The **no** form of the command removes the mirroring session and any mirroring source previously assigned to that session. To preserve the session while deleting a mirroring source assigned to it, refer to the **no** command descriptions under “4. Configure the Monitored Traffic in a Mirror Session” on page B-55.*

< 1 - 4 >: Identifies the mirroring session created by this command. (Multiple sessions on the switch can use the same exit port.)

name < name-str >: Optional alphanumeric name string used to identify the session (up to 15 characters in length).

port < exit-port-# >: Exit port for mirrored traffic in the remote session. This is the port to which a traffic analyzer or IDS is connected.

Configuring a Source Switch in a Remote Mirroring Session. For a remote mirroring session, enter the **mirror remote ip** command on the source switch to configure a remote destination switch for a mirroring session on the source switch. The source IP address, UDP port number, and destination IP address, that you enter must be the same values that you entered with the **mirror endpoint ip** command in “2. Configure a Mirroring Destination on a Remote Switch” on page B-50.

Caution

After you configure a mirroring session with traffic-selection criteria and a destination, the switch immediately starts to mirror traffic to the destination device connected to each exit port. In a remote mirroring session which uses IPv4 encapsulation, if the remote (endpoint) switch is not already configured as the destination for the session, its performance may be adversely affected by the stream of mirrored traffic. For this reason, ProCurve strongly recommends that you configure the endpoint switch in a remote mirroring session

as described in “2. Configure a Mirroring Destination on a Remote Switch” on page B-50, before using the **mirror remote ip** command in this section to configure the mirroring source for the same session.

Syntax: [no] mirror < 1 - 4 > [name < name-str >] remote ip < src-ip >
< src-udp-port > < dst-ip >

This command is used on the source switch to uniquely associate the mirrored traffic in the specified session with a remote destination switch. You must configure the same source and destination parameters when you configure the same session on both the source and destination switches. (If multiple remote sessions use the same source and destination IP addresses, each session must use a unique UDP port value.)

When you execute this command, the following message is displayed:

Caution: Please configure destination switch first.

Do you want to continue [y/n]?

- *If you have not yet configured the session on the remote destination switch, follow the configuration procedure in Step 2 on page B-50 before using this command.*
- *If you have already configured the session on the remote destination switch, type **y** (for “yes”) to complete this command.*

*The **no** form of the command removes the mirroring session and any mirroring source previously assigned to the session. To preserve the session while deleting a monitored source assigned to it, refer to the **no** command descriptions in “4. Configure the Monitored Traffic in a Mirror Session” on page B-55.*

< 1 - 4 >: *Identifies the mirroring session created by this command.*

name < name-str >: *Optional alphanumeric name string used as an additional session identifier (up to 15 characters in length).*

< src-ip >: *The IP address of the VLAN or subnet on which the traffic to be mirrored enters or leaves the switch.*

Syntax: [no] mirror < 1 - 4 > [name < name-str >] remote ip < src-ip >
< src-udp-port > < dst-ip >

< src-udp-port >: *This parameter associates the remote session with a UDP port number. When multiple sessions have the same source IP address < src-ip > and destination IP address < dst-ip >, the UDP port number must be unique in each session. The UDP port number used for a given session should be in the range of 7933 - 65535.*

Caution: *UDP port numbers below 7933 are reserved for various IP applications. Using them for mirroring can result in the interruption of other IP functions, and in non-mirrored traffic being received on the destination switch and sent to a device connected to the remote exit port.*

*The configured UDP port number is included in the frames mirrored from the source switch to the remote destination switch (**mirror endpoint**), and enables the remote switch to match the frames to the exit port configured for the combined UDP port number, source IP address, and destination IP address. Refer to the **mirror endpoint ip** command syntax in “2. Configure a Mirroring Destination on a Remote Switch” on page B-50 for more information.*

< dst-ip >: *For the remote session specified in the command, this is the IP address of the VLAN or subnet on which the remote exit port exists. (The exit port to which a traffic analyzer or IDS is connected is configured on the remote switch in Step 2; see “2. Configure a Mirroring Destination on a Remote Switch” on page B-50.)*

4. Configure the Monitored Traffic in a Mirror Session

This step configures one or more interfaces on a source switch with traffic-selection criteria to select the traffic to be mirrored in a local or remote session configured in Step 3.

Traffic Selection Options

To configure traffic mirroring, specify the source interface, traffic direction, and criteria to be used to select the traffic to be mirrored by using the following options:

- Interface type
 - Port, trunk, and/or mesh
 - VLAN
 - Switch (global configuration level)
- Traffic direction and selection criteria
 - All inbound and/or outbound traffic on a port or VLAN interface
 - Only inbound IP traffic selected with an ACL (deprecated in software release **K.14.01 and greater**)
 - Only inbound IPv4 or IPv6 traffic selected with a classifier-based mirroring policy
 - All inbound and/or outbound traffic selected by MAC source and/or destination address

The different ways to configure traffic-selection criteria on a monitored interface are described in the following sections:

- “Selecting All Inbound/Outbound Traffic to Mirror” on page B-57
- “Selecting Inbound Traffic Using an ACL (Deprecated)” on page B-62
- “Selecting Inbound/Outbound Traffic Using a MAC Address” on page B-63
- “Selecting Inbound Traffic Using Advanced Classifier-Based Mirroring” on page B-66

Mirroring-Source Restrictions

In a mirroring session, you can configure any of the following sources of mirrored traffic:

- Multiple port and trunk, and/or mesh interfaces
- One VLAN

If you configure a VLAN as the source interface in a mirroring session and assign a second VLAN to the session, the second VLAN overwrites the first VLAN as the source of mirrored traffic.

- One classifier-based policy

If you configure a mirroring policy on a port or VLAN interface to mirror inbound traffic in a session, you cannot configure a port, trunk, mesh, ACL, or VLAN as an additional source of mirrored traffic in the session.

- Up to 320 MAC addresses (used to select traffic according to source and/or destination MAC address) in all mirroring sessions configured on a switch

Selecting All Inbound/Outbound Traffic to Mirror

Use the commands in this section to configure all inbound and/or outbound traffic on specified VLAN, port, or trunk interfaces for a local or remote mirroring session. For an example of a mirroring configuration that selects all inbound or outbound traffic on a monitored interface, see:

- “Example: Local Mirroring Using Traffic-Direction Criteria” on page B-87
- “Example: Remote Mirroring Using Traffic-Direction Criteria” on page B-90

Note

If you have already configured session 1 with a local or remote destination (as described in “3. Configure a Mirroring Session on the Source Switch” on page B-52), you can enter the **vlan < vid > monitor** or **interface < port > monitor** command without additional parameters for traffic-selection criteria and session number to configure mirroring for all inbound and outbound traffic on the specified VLAN or port interfaces in session 1 with the preconfigured destination.

Port Interface with Traffic Direction as the Selection Criteria

Use the following command to select all traffic on a port, trunk, and/or mesh interface for mirroring according to traffic direction (inbound and/or outbound):

Syntax: [no] interface < port/trunk/mesh > monitor all < in | out | both > mirror
 < 1 - 4 | name-str > [< 1 - 4 | name-str > < 1 - 4 | name-str >
 < 1 - 4 | name-str >] [no-tag-added]

This command assigns a mirroring source to a previously configured mirroring session on a source switch by specifying the port, trunk, and/or mesh source(s) to use, the direction of traffic to mirror, and the session.

*The **no** form of the command removes a mirroring source assigned to the session, but does not remove the session itself. This enables you to repurpose a session by removing an unwanted mirroring source and adding another in its place.*

interface < port/trunk/mesh >: Identifies the source port(s), static trunk(s), and/or mesh on which to mirroring traffic. Use a hyphen for a range of consecutive ports or trunks (a5-a8, Trk2-Trk4). Use a comma to separate non-contiguous interfaces (b11, b14, Trk4, Trk7).

monitor all < in | out | both >: For the interface specified by **< port/trunk/mesh >**, selects traffic to mirror based on whether the traffic is entering or leaving the switch on the interface.

in: Mirrors entering traffic.

out: Mirrors exiting traffic.

both: Mirrors traffic entering and exiting.

If you enter the **monitor all** command without selection criteria or a session identifier, the command applies by default to session 1 (see “Monitor Command” on page B-96).

mirror < 1 - 4 | < name-str >: Assigns the traffic specified by the interface and direction to a session by number or (if configured) by name. The session must have been previously configured as described in “3. Configure a Mirroring Session on the Source Switch” on page B-52. Depending on how many sessions are already configured on the switch, you can use the same command to assign the specified source to up to four sessions; for example, **interface a1 monitor all in mirror 1 2 4**. For limits on configuring mirroring sources in a session, refer to “Mirroring-Source Restrictions” on page B-56.

< 1 - 4 >: Configures the selected port traffic to be mirrored in the specified session number.

[name < name-str >]: Optional; configures the selected port traffic to be mirrored in the specified session name. The string can be used interchangeably with the session number when using this command to assign a mirroring source to a session. To configure an alphanumeric name for a mirror session, refer to the command description in “Configuring a Source Switch in a Remote Mirroring Session” on page B-53.

[no-tag-added]: Prevents a VLAN tag from being added to the mirrored copy of an outbound packet sent to a local or remote mirroring destination. See “Untagged Mirrored Packets” on page B-59 for more information.

Untagged Mirrored Packets

Although a VLAN tag is added (by default) to the mirrored copy of untagged outbound packets to indicate the source VLAN of the packet, it is sometimes desirable to have mirrored packets look exactly like the original packet. The **no-tag-added** parameter gives you the option of not tagging mirrored copies of outbound packets.

```
ProCurve(config)#interface 3 monitor all in mirror 1 no-tag-added
ProCurve(config)#interface mesh monitor all both mirror 1 no-tag-added
```

Figure B-28. Mirroring Commands with the no-tag-added Option

```
ProCurve# show monitor 1
Network Monitoring
  Session: 1   Session Name:
  ACL: no ACL relationship exists
  Mirror Destination: 48
  Untagged traffic : untagged ← Indicates the no-tag-added option is configured.
  Monitoring Sources Direction
  -----
  Port: 3           Both
```

Figure B-29. Displaying a Mirror Session Configuration with the no-tag-added Option

Using SNMP to Configure No-Tag-Added. The MIB object `hpicfBridgeDontTagWithVlan` is used to implement the no-tag-added option, as shown below:

```
hpicfBridgeDontTagWithVlan OBJECT-TYPE
    SYNTAX INTEGER
        {
            enabled(1),
            disabled(2)
        }
    MAX-ACCESS read-write
    STATUS current
    DESCRIPTION
        "This oid mentions whether VLAN tag is part of the
        mirror'ed copy of the packet. The value 'enabled'
        denotes that the VLAN tag shouldn't be part
```

```
of the mirror'ed copy; 'disabled' does put
the VLAN tag in the mirror'ed copy. Only one
logical port is allowed.
This object is persistent and when written
the entity
SHOULD save the change to non-volatile storage."
DEFVAL { 2 }
::= { hpicfBridgeMirrorSessionEntry 2 }
```

Operating Notes. The following conditions apply for the **no-tag-added** option:

- The specified port can be a physical port, trunk port, or mesh port.
- Only a single logical port (physical port or trunk) can be associated with a mirror session when the **no-tag-added** option is specified. No other combination of ACL mirroring, VLAN mirroring, or port mirroring can be associated with the mirror session. If more than one logical port is specified, the following error message is displayed:

Cannot monitor more than one logical port with no-tag-added option

- If a port changes its VLAN membership and/or untagged status within the VLAN, the “untagged port mirroring” associated with that port is updated when the configuration change is processed.
- Only four ports or trunks can be monitored at one time when all four mirror sessions are in use (one logical port per mirror session) without VLAN tags being added to a mirrored copy.
- The **no-tag-added** option can also be used when mirroring is configured with SNMP.
- A VLAN tag is still added to the copies of untagged packets obtained via VLAN-based mirroring.

VLAN Interface with Traffic Direction as the Selection Criteria

Use the following command to select all traffic on a VLAN interface for mirroring according to traffic direction (inbound and/or outbound):

Syntax: `vlan < vid-#> monitor all < in | out | both > mirror < 1 - 4 | name-str > [< 1 - 4 | name-str > < 1 - 4 | name-str > < 1 - 4 | name-str >]`

This command assigns a monitored VLAN source to a previously configured mirroring session on a source switch by specifying the VLAN ID, the direction of traffic to mirror, and the session. Assigning a VLAN to a mirroring session precludes assigning any other mirroring sources to the same session. If a VLAN is already assigned to a given mirroring session, using this command to assign another VLAN to the same mirroring session results in the second assignment replacing the first. Also, if there are other (port, trunk, or mesh) mirroring sources already assigned to a session, the switch displays a message similar to:

`Mirror source port exists on session N. Can not add mirror source VLAN.`

*The **no** form of the command removes a mirroring source assigned to the session, but does not remove the session itself. This allows you to repurpose a session by removing an unwanted mirroring source and adding another in its place.*

vlan < vid-#>: *Identifies the VLAN on which to mirror traffic.*

monitor all < in | out | both >: *Uses the direction of traffic on the specified vid-# to select traffic to mirror. Refer to the syntax description on page B-57. (If you enter the **monitor all** command without selection criteria or a session identifier, the command applies by default to session 1; see “Monitor Command” on page B-96.)*

mirror < 1 - 4 | < name-str >: *Assigns the VLAN traffic defined by the VLAN ID and traffic direction to a session number or name. (The session must have been previously configured as described in “3. Configure a Mirroring Session on the Source Switch” on page B-52.) Depending on how many sessions are already configured on the switch, you can use the same command to assign the specified VLAN source to up to four sessions; for example, **interface a1 monitor all in mirror 1 2 4**. For limits on configuring mirroring sources in a session, refer to “Mirroring-Source Restrictions” on page B-56.*

< 1 - 4 >: Configures the selected VLAN traffic to be mirrored in the specified session number.

[name < name-str >]: Optional; configures the selected port traffic to be mirrored in the specified session name. The string can be used interchangeably with the session number when using this command to assign a mirroring source to a session. To configure an alphanumeric name for a mirroring session, refer to the command description under “Configuring a Source Switch in a Remote Mirroring Session” on page B-53.

Selecting Inbound Traffic Using an ACL (Deprecated)

Deprecation of ACL-based Traffic Selection

In release K.14.01 or greater, the use of ACLs to select inbound traffic in a mirroring session has been replaced with classifier-based mirroring policies (see “Selecting Inbound Traffic Using Advanced Classifier-Based Mirroring” on page B-66).

The following commands have been deprecated:

- **interface <port/trunk/mesh> monitor ip access-group <acl-name> in mirror < 1 - 4 | name-str >**
- **vlan < vid-# > monitor ip access-group <acl-name> in mirror < 1 - 4 | name-str >**

After you install and boot release K.14.01 or greater, ACL-based local and remote mirroring sessions configured on a port or VLAN interface are automatically converted to classifier-based mirroring policies. For more information, see “Migration to Release K.14.01 or Greater” on page B-37.

Selecting Inbound/Outbound Traffic Using a MAC Address

Use the **monitor mac mirror** command at the global configuration level to apply a source and/or destination MAC address as the selection criteria used in a local or remote mirroring session.

While classifier-based mirroring allows you to mirror traffic using a policy to specify IP addresses as selection criteria, MAC-based mirroring allows you monitor switch traffic using a source and/or destination MAC address. You can apply MAC-based mirroring in one or more mirroring sessions on the switch to monitor:

- Inbound traffic
- Outbound traffic
- Both inbound and outbound traffic

MAC-based mirroring is useful in ProCurve Network Immunity security solutions that provide detection and response to malicious traffic at the network edge. After isolating a malicious MAC address, a security administrator can mirror all traffic sent to, and received from, the suspicious address for troubleshooting and traffic analysis.

The MAC address that you enter with the **monitor mac mirror** command is configured to select traffic for mirroring from all ports and learned VLANs on the switch. Therefore, a suspicious MAC address used in wireless applications can be continuously monitored as it re-appears in switch traffic on different ports or VLAN interfaces.

You can configure MAC-based mirroring from the CLI or an SNMP management station and use it to mirror:

- All inbound and outbound traffic from a group of hosts to one destination device.
- Inbound and/or outbound traffic from each host to a different destination device.
- Inbound and outbound traffic from all monitored hosts separately on two destination devices: mirroring all inbound traffic to one device and all outbound traffic to another device.

To configure a MAC address to filter mirrored traffic on an interface, enter the **monitor mac mirror** command at the global configuration level.

Syntax: [no] monitor mac <mac-addr> <src | dest | both> mirror <1 - 4 | name-str > [<1 - 4 | name-str >] [<1 - 4 | name-str >] [<1 - 4 | name-str >]

*Use this command to configure a source and/or destination MAC address as criteria for selecting traffic in one or more mirroring sessions on the switch. The MAC address you enter is configured to mirror inbound (**src**), outbound (**dest**), or both inbound and outbound (**both**) traffic on any port or learned VLAN on the switch.*

Packets that are sent or received on an interface configured with a mirroring session and contain the MAC address as source and/or destination address are mirrored to a previously configured destination device.

*To remove a MAC address as selection criteria in a mirroring session, you must enter the complete command syntax; for example, **no monitor mac 998877-665544 dest mirror 4**.*

*The **no** form of the command removes the MAC address as a mirroring criteria from an active session, but does not remove the session itself. This enables you to repurpose a session by removing an unwanted mirroring criteria and adding another in its place.*

monitor mac < mac-addr>: Configures the MAC address as selection criteria for mirroring traffic on any port or learned VLAN on the switch.

< src | dest | both >: Specifies how the MAC address is used to filter and mirror packets in inbound and/or outbound traffic on the interfaces on which the mirroring session is applied:

src: Mirrors all packets in inbound traffic that contain the specified MAC address as source address.

dest: Mirrors all packets in outbound traffic that contain the specified MAC address as destination address.

Note: The MAC address of the switch is not supported as either the source or destination MAC address used to select mirrored traffic.

both: Mirrors all packets in both inbound and outbound traffic that contain the specified MAC address as either source or destination address.

mirror < 1 - 4 | < name-str >: Assigns the inbound and/or outbound traffic filtered by the specified MAC address to a previously configured mirroring session. The session is identified by a number or (if configured) a name.

*Depending on how many sessions are configured on the switch, you can use the same command to configure a MAC address as mirroring criteria in up to four sessions. To identify a session, you can enter either its name or number; for example: **mirror 1 2 3 traffsrc4***

Refer to “Mirroring-Source Restrictions” on page B-56 for the restrictions on how many mirroring source criteria you can configure in the same session.

< 1 - 4 >: Specifies a mirroring session by number (1 to 4), for which the configured MAC address is used to select and mirror inbound and/or outbound traffic.

[name < name-str >]: (Optional) Specifies a mirroring session by name (alphanumeric string), for which the configured MAC address is used to select and mirror inbound and/or outbound traffic. For a remote mirroring session, you must configure the same session name on both the source and destination switch.

Restrictions

The following restrictions apply to MAC-based mirroring:

- Up to 320 different MAC addresses are supported for traffic selection in all mirroring sessions configured on the switch.
- A destination MAC address is not supported as mirroring criteria for routed traffic because in routed packets, the destination MAC address is changed to the next-hop address when the packet is forwarded. Therefore, the destination MAC address that you want to mirror will not appear in routed packet headers.

This restriction also applies to the destination MAC address of a host that is directly connected to a routing switch. (Normally, a host is connected to an edge switch, which is directly connected to the router.)

To mirror routed traffic, it is recommended that you use classifier-based policies to select IPv4 or IPv6 traffic for mirroring as described in “Selecting Inbound Traffic Using Advanced Classifier-Based Mirroring” on page B-66.

- On a switch, you can use a MAC address only once as a source MAC address, and only once as a destination MAC address, to filter mirrored traffic.

For example, after you enter the following commands:

```
monitor mac 111111-222222 src mirror 1  
monitor mac 111111-222222 dest mirror 2
```

The following commands are not supported:

```
monitor mac 111111-222222 src mirror 3  
monitor mac 111111-222222 dest mirror 4
```

In addition, if you enter the **monitor mac 111111-222222 both mirror 1** command, you cannot use the MAC address **111111-222222** in any other **monitor mac mirror** configuration commands on the switch.

- To re-use a MAC address that has already been configured as a source and/or destination address for traffic selection in a mirror session, you must first remove the configuration by entering the **no** form of the command, and then re-enter the MAC address in a new **monitor mac mirror** command.

For example, if you have already configured MAC address **111111-222222** to filter inbound and outbound mirrored traffic, and decide to use it to filter only inbound traffic in a mirror session, you could enter the following commands:

```
monitor mac 111111-222222 both mirror 1  
no monitor mac 111111-222222 both mirror 1  
monitor mac 111111-222222 src mirror 1
```

- A mirroring session in which you configure MAC-based mirroring is not supported on a port, trunk, mesh or VLAN interface on which a mirroring session with a classifier-based mirroring policy is configured.

Selecting Inbound Traffic Using Advanced Classifier-Based Mirroring

In software release K.14.01 or greater, in addition to the traffic selection options described in “4. Configure the Monitored Traffic in a Mirror Session” on page B-55, traffic mirroring supports the use of advanced classifier-based functions that provide:

- A finer granularity for selecting the inbound IP traffic that you want to mirror on an individual port or VLAN interface (instead of mirroring all inbound traffic on the interface)
- Support for mirroring both IPv4 and IPv6 traffic
- The ability to re-use the same traffic classes in different software-feature configurations; for example, you can apply both a QoS rate-limiting and mirroring policy on the same class of traffic.

**Deprecation of
ACL-based
Traffic
Selection**

In software release K.14.01 or greater, advanced classifier-based policies replace ACL-based traffic selection in mirroring configurations.

Like ACL-based traffic-selection criteria, classifier-based service policies apply only to inbound traffic flows and are configured on a per-port or per-VLAN basis. In a mirroring session, classifier-based service policies do not support:

- The mirroring of outbound traffic exiting the switch
- The use of meshed ports as monitored (source) interfaces

Classifier-based mirroring is *not* designed to work with other traffic-selection methods in a mirroring session applied to a port or VLAN interface:

- If a mirroring session is already configured with one or more traffic-selection criteria (MAC-based or all inbound and/or outbound traffic), the session does not support the addition of a classifier-based policy.
- If a mirroring session is configured to use a classifier-based mirroring policy, no other traffic-selection criteria (MAC-based or all inbound and/or outbound traffic) can be added to the session on the same or a different interface.

Classifier-based mirroring policies provide greater precision when analyzing and debugging a network traffic problem. Using multiple match criteria, you can finely select and define the classes of traffic that you want to mirror on a traffic analyzer or IDS device.

For more information on how to configure and use classifier-based service policies, refer to the “Classifier-Based Software Configuration” chapter in the *Advanced Traffic Management Guide*.

For an example of a mirroring configuration that uses a classifier-based service policy to select traffic on a monitored interface, see “Example: Remote Mirroring Using a Classifier-Based Policy” on page B-88.

Classifier-Based Mirroring Configuration

To use the classifier-based model to configure a mirroring policy and apply it to a selected class of traffic on a port or VLAN interface, follow these steps:

1. Evaluate the types of traffic in your network and identify the traffic types that you want to mirror.

2. Create an IPv4 or IPv6 traffic class using the **class** command to select the packets that you want to mirror in a session on a preconfigured local or remote destination device.

Context: Global configuration

Syntax: [no] class < ipv4 | ipv6 > <classname >

Defines the name of a traffic class and specifies whether a policy is to be applied to IPv4 or IPv6 packets, where < classname > is a text string (64 characters maximum).

*After you enter the **class** command, you enter the class configuration context to specify match criteria. A traffic class contains a series of **match** and **ignore** commands, which specify the criteria used to classify packets.*

*To configure a default traffic class, use the **default-class** command as described below. A default class manages the packets that do not match the match/ignore criteria in any other classes in a policy.*

A traffic class consists of match criteria, which consist of **match** and **ignore** commands.

- **match** commands define the values that header fields must contain for a packet to belong to the class and be managed by policy actions.
- **ignore** commands define the values which, if contained in header fields, exclude a packet from the policy actions configured for the class.

Note

Be sure to enter match/ignore statements in the *precise order* in which you want their criteria to be used to check packets.

The following match criteria are supported in match/ignore statements for inbound IPv4/IPv6 traffic:

- IP source address (IPv4 and IPv6)
- IP destination address (IPv4 and IPv6)
- IP protocol (such as ICMP or SNMP)
- Layer 3 IP precedence bits
- Layer 3 DSCP codepoint
- Layer 4 TCP/UDP application port (including TCP flags)
- VLAN ID

Enter one or more **match** or **ignore** commands from the class configuration context to filter traffic and determine the packets on which policy actions will be performed.

Context: Class configuration

Syntax: [no] [seq-number] < match | ignore > < ip-protocol >
< source-address > < destination-address > [ip-dscp codepoint]
[precedence precedence-value] [tos tos-value] [vlan vlan-id]

For detailed information about how to enter **match** and **ignore** commands to configure a traffic class, refer to the “Creating a Traffic Class” section in the “Classifier-Based Software Configuration” in the *Advanced Traffic Management Guide*.

3. Create a mirroring policy to configure the session and destination device to which specified classes of inbound traffic are sent by entering the **policy mirror** command from the global configuration context.

Context: Global configuration

Syntax: [no] policy mirror <policy-name >

Defines the name of a mirroring policy and enters the policy configuration context.

A traffic policy consists of one or more classes, and one or more mirroring actions configured for each class of traffic. The configured actions are executed on packets that match a **match** statement in a class. No policy action is performed on packets that match an **ignore** statement.

Note

Be sure to enter each class and its associated mirroring actions in the *precise order* in which you want packets to be checked and processed.

To configure the mirroring actions that you want to execute on packets that match the criteria in a specified class, enter one or more **class action mirror** commands from the policy configuration context:

Context: Policy configuration

Syntax: [no] [seq-number] class < ipv4 | ipv6 > <classname >
action mirror <session >

*Defines the mirroring action to be applied on a pre-configured IPv4 or IPv6 traffic class when a packet matches the **match** criteria in the traffic class. You can enter multiple **class action mirror** statements in a policy. You can configure only one mirroring session (destination) for each class. You can configure the same mirroring session for different classes.*

- **[seq-number]** — The (optional) **seq-number** parameter sequentially orders the mirroring actions that you enter in a policy configuration. Actions are executed on matching packets in numerical order. Default: Mirroring action statements are numbered in increments of 10, starting at 10.
- **class < ipv4 | ipv6 > <classname >** — Defines the preconfigured traffic class on which the mirroring actions in the policy are executed, and specifies whether the mirroring policy is applied to IPv4 or IPv6 traffic in the class. The classname is a text string (64 characters maximum).
- **action mirror <session >** — Configures mirroring for the destination and session specified by the **session** parameter.

A packet that matches the **match** criteria in a class is mirrored to the exit (local or remote) port that has been previously configured for the session, where **session** is a value from **1** to **4** or a text string (if you configured the session with a name when you entered the **mirror** command).

Prerequisite: The local or remote exit port for a session must already be configured before you enter the **mirror < session >** parameter in a **class action** statement:

- In a local mirroring session, the exit port is configured with the **mirror <session-number> port** command.
- In a remote mirroring session, the remote exit port is configured with the **mirror endpoint ip** and **mirror <session-number> remote ip** commands.

See “2. Configure a Mirroring Destination on a Remote Switch” on page B-50 and “3. Configure a Mirroring Session on the Source Switch” on page B-52 for more information.

Restriction: In a policy, you can configure only one mirroring session per class. You can configure the same session for different classes.

Mirroring is not executed on packets that match **ignore** criteria in a class.

The execution of mirroring actions is performed in the order in which the classes are numerically listed in the policy.

The complete **no** form of the **class action mirror** command or the **no <seq-number >** command removes a class and mirroring action from the policy configuration.

To manage packets that do not match the **match** or **ignore** criteria in any class in the policy, and therefore have no mirroring actions performed on them, you can enter an optional default class. The default class is placed at the end of a policy configuration and specifies the mirroring actions to perform on packets that are neither matched nor ignored.

4. (Optional) To configure a default-class in a policy, enter the **default-class** command at the end of a policy configuration and specify one or more actions to be executed on packets that are not matched and not ignored.

Context: Policy configuration

Syntax: [no] default-class action mirror <session> [action mirror <session> ...]

Configures a default class that allows packets that are not matched nor ignored by any of the class configurations in a mirroring policy to be mirrored to the destination configured for the specified session.

Prerequisite: *The local or remote exit port for a session must already be configured with a destination device before you enter the **mirror <session>** parameter in a **default-class action** statement. See “2. Configure a Mirroring Destination on a Remote Switch” on page B-50 and “3. Configure a Mirroring Session on the Source Switch” on page B-52 for more information.*

For general information about how to configure and manage a service policy, refer to the “Creating a Service Policy” section in the “Classifier-Based Software Configuration” chapter in the *Advanced Traffic Management Guide*.

5. Apply the mirroring policy to inbound traffic on a port (**interface service-policy in** command) or VLAN (**vlan service-policy in** command) interface.

Caution

After you apply a mirroring policy for one or more preconfigured sessions on a port or VLAN interface, the switch immediately starts to use the traffic-selection criteria and exit port to mirror traffic to the destination device connected to each exit port.

In a remote mirroring session which uses IPv4 encapsulation, if the remote switch is not already configured as the destination for the session, its performance may be adversely affected by the stream of mirrored traffic.

For this reason, ProCurve strongly recommends that you first configure the exit switch in a remote mirroring session, as described in “2. Configure a Mirroring Destination on a Remote Switch” on page B-50 and “3. Configure a Mirroring Session on the Source Switch” on page B-52, before you apply a mirroring service policy on a port or VLAN interface.

The following restrictions apply to a mirroring service policy:

- Only one mirroring policy is supported on a port or VLAN interface.
- If you apply a mirroring policy to a port or VLAN interface on which a mirroring policy is already configured, an error message is displayed. The new policy does not overwrite the existing one. To apply a new policy, you must first remove the existing policy with the **no interface service-policy in** or **no vlan service-policy in** command.
- A mirroring policy is supported only on inbound traffic.

Because only one mirroring policy is supported on a port or VLAN interface, ensure that the policy you want to apply contains all the required classes and actions for your configuration.

To apply a mirroring policy on a port or VLAN interface, enter one of the following **service-policy** commands from the global configuration context.

Context: Global configuration

Syntax: interface <port-list> service-policy <policy-name> in

Configures the specified port(s) with a mirroring policy that is applied to inbound traffic on each interface.

*Separate individual port numbers in a series with a comma; for example, **a1, b4, d3**. Enter a range of ports by using a dash; for example, **a1-a5**.*

*The mirroring policy name you enter must be the same as the policy name you configured with the **policy mirror** command in Step 2.*

Syntax: vlan <vlan-id> service-policy <policy-name> in

Configures a mirroring policy on the specified VLAN that is applied to inbound traffic on the VLAN interface.

Valid VLAN ID numbers range from 1 to 4094.

*The mirroring policy name you enter must be the same as the policy name you configured with the **policy mirror** command in Step 2*

For more information about how to apply a mirroring policy to an interface, refer to the “Applying a Service Policy to an Interface” section in the “Classifier-Based Software Configuration” chapter in the *Advanced Traffic Management Guide*.

Viewing a Classifier-Based Mirroring Configuration

To display information about a classifier-based mirroring configuration or statistics on one or more mirroring policies, enter one of the following commands:

- **show class** < *class-name* >
- **show policy** < *mirror-policy-name* >
- **show policy resources**
- **show statistics policy** [*mirror-policy-name*] [**interface** <*port-list*> | **vlan** <*vlan-id*>] **in**

For examples of classifier-based show command output, see “Displaying Information on a Classifier-Based Mirroring Session” on page B-82.

Classifier-Based Mirroring Restrictions

The following restrictions apply to mirroring policies configured with the classifier-based model:

- A mirroring policy is supported only on *inbound* IPv4 or IPv6 traffic.
- A mirroring policy is not supported on a meshed port interface. (Classifier-based policies are supported only on a port, VLAN, or trunk interface.)
- Only one classifier-based mirroring policy is supported on a port or VLAN interface. You can, however, apply a classifier-based policy of a different type, such as QoS.
- You can enter multiple **class action mirror** statements in a policy.
 - You can configure only one mirroring session (destination) for each class.
 - You can configure the same mirroring session for different classes.

- If a mirroring session is configured with a classifier-based mirroring policy on a port or VLAN interface, no other traffic-selection criteria (MAC-based or all inbound and/or outbound traffic) can be added to the session.

```
Switch-B(config)# mirror endpoint 10.10.40.4 9200 10.10.50.5 port a1
...
Switch-A(config)# mirror 1 remote ip 10.10.40.4 9200 10.10.50.5
Caution: Please configure destination switch first.
Do you want to continue [y/n]? y
Switch-A(config)# class ipv4 Data2
Switch-A(config-class)# match ip 10.28.31.1 any
Switch-A(config-class)# match ip any host 10.28.31.0/24
Switch-A(config-class)# exit
Switch-A(config)# policy mirror SalesData
Switch-A(config-policy)# class ipv4 Data2 action mirror 1
Switch-A(config-policy)# exit
Switch-A(config)# vlan 10 service-policy SalesData in
Switch-A(config)# vlan 10 monitor all out mirror 1
A prior mirror policy relationship exists with mirror session 1. Please remove.
```

Classifier-based policy used to select mirrored traffic in session 1

The configuration of additional traffic-direction criteria to select mirrored traffic is not supported in session 1.

Figure B-30. Mirroring Configuration in Which Only a Mirroring Policy is Supported

- If a mirroring session is already configured with one or more traffic-selection criteria (MAC-based or all inbound and/or outbound traffic), the session does not support the addition of a classifier-based policy.

```
Switch-B(config)# mirror endpoint 10.10.40.4 9200 10.10.50.5 port a1
...
Switch-A(config)# mirror 1 remote ip 10.10.40.4 9200 10.10.50.5
Caution: Please configure destination switch first.
Do you want to continue [y/n]? y
Switch-A(config)# vlan 10 monitor all out mirror 1
Switch-A(config)# class ipv4 Data2
Switch-A(config-class)# match ip 10.28.31.1 any
Switch-A(config-class)# match ip any host 10.28.31.0/24
Switch-A(config-class)# exit
Switch-A(config)# policy mirror SalesData
Switch-A(config-policy)# class ipv4 Data2 action mirror 1
Switch-A(config-policy)# exit
Switch-A(config)# vlan 10 service-policy SalesData in
Mirror source VLAN exists on mirror session 1. Cannot add this mirror source.
```

Configuration of traffic-direction criteria to select all outbound traffic on VLAN 10 in mirror session 1

The configuration of an additional classifier-based policy to select mirrored traffic on VLAN 10 is not supported in session 1.

Figure B-31. Mirroring Configuration in Which Only Traffic-Selection Criteria are Supported

Applying Multiple Mirroring Sessions to an Interface

You can apply a mirroring policy to an interface that is already configured with another traffic-selection method (MAC-based or all inbound and/or outbound traffic) for a different mirroring session.

The classifier-based policy provides a finer level of granularity that allows you to zoom in on a subset of port or VLAN traffic and select it for local or remote mirroring.

In the following example, traffic on Port b1 is used as the mirroring source for two different, local mirroring sessions:

- All inbound and outbound traffic on Ports b1, b2, and b3 is mirrored in session 4.
- Only selected voice traffic on Port b1 is mirrored in session 2.

```
ProCurve(config)# mirror 4 port a2
ProCurve(config)# interface b1-b3 monitor all both mirror 4
ProCurve(config)# mirror 2 port b4
ProCurve(config)# class ipv4 voice
ProCurve(config-class)# match ip any any ip-dscp ef
ProCurve(config-class)# exit
ProCurve(config)# policy mirror IPphones
ProCurve(config-policy)# class ipv4 voice action mirror 2
ProCurve(config-policy)# exit
ProCurve(config)# interface b1 service-policy IPphones in
```

Figure B-32. Example of Applying Multiple Sessions to the Same Interface

Displaying a Mirroring Configuration

Displaying All Mirroring Sessions Configured on the Switch

Use the **show monitor** command to display information on the currently configured status, traffic-selection criteria, and number of monitored interfaces in each mirroring session on a switch. The exit ports configured on the switch for remote mirroring sessions (remote endpoints) are also displayed.

```
ProCurve# show monitor
```

Sessions	Status	Type	Sources	Policy
1	active	port	1	yes
2	active	mac	2	no
3	not defined			
4	inactive	IPv4	0	no

Network Monitoring

Remote Mirroring - Remote Endpoints

Type	UDP Source Addr	UDP port	UDP Dest Addr	Dest Port
IPv4	10.10.30.1	7950	10.10.20.1	B10

Local and Remote Mirroring Sources:

- **Session 1** is performing local mirroring using a classifier-based policy as traffic-selection criteria.
- **Session 2** is performing remote mirroring using MAC-based traffic-selection criteria.
- **Session 3** is not configured.
- **Session 4** is configured for remote mirroring from a non-policy source (for example, traffic direction), but is currently not mirroring any traffic.

Remote Mirroring Destination:

The switch is configured as a remote mirroring destination (endpoint) for a source at 10.10.30.1, using port B10 as the exit port.

Figure B-33. Displaying the Currently Configured Mirroring Sessions on the Switch

Syntax: show monitor

*If a monitored source for a remote session is configured on the switch, the following information is displayed. Otherwise, the output displays: **Mirroring is currently disabled.***

Sessions: Lists the four configurable sessions on the switch.

Status: Displays the current status of each session:

active: The session is configured.

inactive: The session is partially configured. Only the destination has been configured; the mirroring source is not configured.

not defined: Mirroring is not configured for this session.

Syntax: show monitor

Type: *Indicates whether the mirroring session is local (**port**), remote (**IPv4**), or MAC-based (**mac**) for local or remote sessions.*

Sources: *Indicates how many monitored source interfaces are configured for each mirroring session.*

Policy: *Indicates whether the source is using a classifier-based mirroring policy to select inbound IPv4 or IPv6 traffic for mirroring.*

*If a remote mirroring endpoint is configured on the switch, then the following information is displayed. Otherwise, the output displays: **There are no Remote Mirroring endpoints currently assigned.***

Type: *Indicates whether the mirroring session is local (**port**), remote (**IPv4**), or MAC-based (**mac**) for local or remote sessions.*

UDP Source Addr: *The IP address configured for the source VLAN or subnet on which the monitored source interface exists. In the configuration of a remote session, the same UDP source address must be configured on the source and destination switches.*

UDP port: *The unique UDP port number that identifies a remote session. In the configuration of a remote session, the same UDP port number must be configured on the source and destination switches.*

UDP Dest Addr: *The IP address configured for the destination VLAN or subnet on which the remote exit port exists. In the configuration of a remote session, the same UDP destination address must be configured on the source and destination switches.*

Dest Port: *Identifies the exit port for a remote session on a remote destination switch.*

Displaying the Remote Endpoints Configured on the Switch

Syntax: show monitor endpoint

*This command displays the remote mirroring endpoints configured on the switch. Information on local sessions configured on the switch is not displayed. (To view the configuration of a local session, use the **show monitor** [**< 1-4 | name < name-str >**] command as described on pages B-76 and B-79.)*

Type: Indicates whether the session is a **port** (local) or **IPv4** (remote) mirroring session.

UDP Source Addr: The IP address configured for the source VLAN or subnet on which the monitored source interface exists. In the configuration of a remote session, the same UDP source address must be configured on the source and destination switches.

UDP port: The unique UDP port number that identifies a remote session. In the configuration of a remote session, the same UDP port number must be configured on the source and destination switches.

UDP Dest Addr: The IP address configured for the destination VLAN or subnet on which the remote exit port exists. In the configuration of a remote session, the same UDP destination address must be configured on the source and destination switches.

Dest Port: Identifies the exit port for a remote session on a remote destination switch.

For example, in Figure B-34, the **show monitor endpoint** output shows that the switch is configured as the remote endpoint (destination) for two remote sessions from the same monitored source interface.

```
ProCurve(config)# show monitor endpoint
Remote Mirroring - Remote Endpoints
```

Type	UDP Source Addr	UDP port	UDP Dest Addr	Dest Port
IPv4	10.10.10.1	8001	10.10.30.2	4
IPv4	10.10.10.1	8003	10.10.30.2	5

These two sessions monitor traffic from the same source switch, but use different UDP port numbers.

Figure B-34. Displaying the Configuration of Remote Mirroring Endpoints on the Switch

Displaying the Mirroring Configuration for a Specific Session

Syntax: show monitor < 1 - 4 | name < name-str >

Use this command to display detailed configuration information for a specified local or remote mirroring session on a source switch.

Session: *Displays the number of the specified session.*

Session Name: *Displays the name of the session, if configured.*

Policy: *Indicates whether the source is using a classifier-based mirroring policy to select inbound IPv4 or IPv6 traffic for mirroring.*

Mirroring Destination: *For a local mirroring session, displays the port configured as the exit port on the source switch. For a remote mirroring session, displays IPv4, which indicates mirroring to a remote (endpoint) switch.*

UDP Source Addr: *The IP address configured for the source VLAN or subnet on which the monitored source interface exists. In the configuration of a remote session, the same UDP source address must be configured on the source and destination switches.*

UDP port: *The unique UDP port number that identifies a remote session. In the configuration of a remote session, the same UDP port number must be configured on the source and destination switches.*

UDP Dest Addr: *The IP address configured for the destination VLAN or subnet on which the remote exit port exists. In the configuration of a remote session, the same UDP destination address must be configured on the source and destination switches.*

Status: *For a remote session, displays current session activity:*

active: *The session is configured and is mirroring traffic. A remote path has been discovered to the destination.*

inactive: *The session is configured, but is not currently mirroring traffic. A remote path has not been discovered to the destination.*

not defined: *Mirroring is not configured for this session.*

Monitoring Sources: *For the specified local or remote session, displays the source (port, trunk, or VLAN) interface and the MAC address (if configured) used to select mirrored traffic.*

Syntax: show monitor < 1 - 4 | name < name-str >

Direction: For the selected interface, indicates whether mirrored traffic is entering the switch (**in**), leaving the switch (**out**), or **both**.

Displaying a Remote Mirroring Session. After you configure session 2 for remote mirroring (Figure B-35), you can enter the **show monitor 2** command to verify the configuration (Figure B-36).

```
ProCurve(config)# mirror 2 name test-10 remote ip 10.10.10.1 8010 10.10.30.2
Caution: Please configure destination switch first.
Do you want to continue [y/n]? y
ProCurve(config)# interface b1 monitor all both mirror 2
```

Figure B-35. Configuring a Remote Mirroring Session and Monitored Source

```
ProCurve_8200(config)# show monitor 2
Network Monitoring

Session: 2      Session Name: test-10
Policy: no policy relationship exists

Mirror Destination: IPv4
  UDP Source Addr  UDP port  UDP Dest Addr  Status
-----
  10.10.10.1      8010     10.10.30.2    active

Monitoring Sources  Direction
-----
Port: B1            Both
```

If no monitored (source) interface is configured for a mirroring session, no information is displayed in the Monitoring Sources and Direction columns.

Figure B-36. Displaying the Configuration of a Remote Mirroring Session

Displaying a MAC-based Mirroring Session. After you configure a MAC-based mirroring session (Figure B-37), you can enter the **show monitor 3** command to verify the configuration (Figure B-38).

```
ProCurve(config)# mirror 3 port a1
ProCurve# monitor mac 112233-445566 src mirror 3
```

Figure B-37. Configuring a MAC-based Mirroring Session

```

ProCurve_8200(config)# show monitor 3
Network Monitoring

Session: 3      Session Name:
Policy: no policy relationship exists

Mirror Destination:  A1      (Port)

Monitoring Sources  Direction
-----
MAC:  112233-445566 Source ←

```

The MAC address used to select packets in a local mirroring session is displayed in these columns.

Figure B-38. Displaying a MAC-based Mirroring Session

Displaying a Local Mirroring Session. When used to display the configuration of a local session, the **show monitor** command displays a subset of the information displayed for a remote mirroring session. For example, Figure B-39 displays a local mirroring configuration for a session configured as follows:

- Session number: 1
- Session name: Detail
- Classifier-based mirroring policy, “MirrorAdminTraffic”, is used to select inbound traffic on port B1.
- Mirrored traffic is sent to exit port B3.

```

ProCurve_8200(config)# show monitor 1
Network Monitoring

Session: 1      Session Name: Detail
Policy: MirrorAdminTraffic

Mirror Destination:  B3      (Port)

Monitoring Sources  Direction
-----
Port: B1           In

```

Figure B-39. Displaying the Configuration of a Local Mirroring Session

Displaying Information on a Classifier-Based Mirroring Session. In the following example, a classifier-based mirroring policy (**mirrorAdminTraffic**) mirrors selected inbound IPv4 packets on VLAN 5 to the destination device configured for mirroring session 3.

```
ProCurve(config)# mirror 3 port c1
Caution: Please configure destination switch first.
Do you want to continue [y/n]? y
ProCurve(config)# class ipv4 AdminTraffic
ProCurve(config-class)# match ip 15.29.61.1 0.63.255.255 0.0.0.0 255.255.255.255
ProCurve(config-class)# match ip 0.0.0.0 255.255.255.255 15.29.61.1 0.63.255.255
ProCurve(config-class)# exit
ProCurve(config)# policy mirror MirrorAdminTraffic
ProCurve(config-policy)# class ipv4 AdminTraffic action mirror 3
ProCurve(config-policy)# exit
ProCurve(config)# vlan 5 service-policy MirrorAdminTraffic in
```

Figure B-40. Configuring a Classifier-Based Mirroring Policy in a Local Mirroring Session

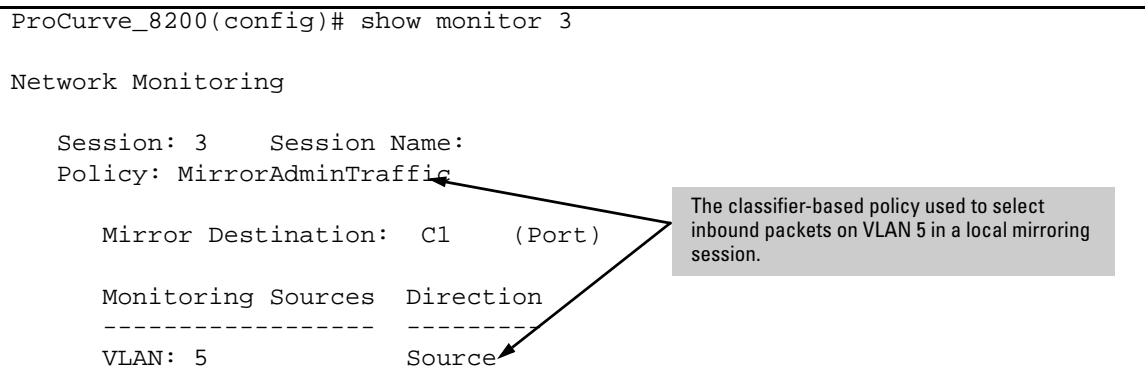
```
ProCurve_8200(config)# show monitor 3

Network Monitoring

Session: 3      Session Name:
Policy: MirrorAdminTraffic

Mirror Destination: C1      (Port)

Monitoring Sources  Direction
-----
VLAN: 5            Source
```



The classifier-based policy used to select inbound packets on VLAN 5 in a local mirroring session.

Figure B-41. Displaying a Classifier-based Policy in a Local Mirroring Session

Use the following **show** commands to display information about:

- A classifier-based mirroring configuration (**show class** and **show policy**)
- Statistics on one or more mirroring policies (**show statistics policy**)
- Hardware resources used by all mirroring policies currently configured on the switch (**show policy resources**).

Syntax: show class < class-name >

Displays the configuration of the specified traffic class, including all match/ignore statements used to classify the packets to be mirrored.

```
ProCurve(config)# show class ipv4 AdminTraffic

Statements for Class ipv4 "AdminTraffic"

10 match ip 15.29.16.1 0.63.255.255 0.0.0.0 255.255.255.255
20 match ip 0.0.0.0 255.255.255.255 15.29.16.1 0.63.255.255
```

Figure B-42. "show class" Output for a Mirroring Policy

Syntax: show policy < policy-name >

Displays the specified policy configuration, including the traffic classes and action commands used to mirror inbound traffic to previously configured destination devices.

```
ProCurve(config)# show policy MirrorAdminTraffic

Statements for Policy "MirrorAdminTraffic"

    10 class ipv4 "AdminTraffic" action mirror 3
```

Figure B-43. "show policy" Output for a Mirroring Policy

Syntax: show statistics policy [*mirror-policy-name*] [interface <*port-list*> | vlan <*vlan-id*>] in

Displays statistics for the specified mirroring policies configured on one or more port or VLAN interfaces.

```
ProCurve# show statistics policy MirrorAdminTraffic vlan 30 in
HitCounts for Policy MirrorAdminTraffic
10 class ipv4 "AdminTraffic" action mirror 3
(5244) 10 match ip 15.29.16.1 0.63.255.255 0.0.0.0 255.255.255.255
(9466) 20 match ip 0.0.0.0 255.255.255.255 15.29.16.1 0.63.255.255
```

Number of packets (in parentheses) that have been mirrored for each match/ignore statement in the mirroring policy

Figure B-44. “show statistics policy” Output for a Mirroring Policy

Displaying Resource Usage for Mirroring Policies

Syntax: show policy resources

Displays the number of hardware resources (rules, meters, and application port ranges) used by classifier-based mirroring policies (local and remote) that are currently applied to interfaces on the switch, as well as QoS policies and other software features.

Note: *The information displayed is the same as the output of the **show qos resources** and **show access-list resources** commands.*

*For a detailed explanation of the information displayed with the **show <qos | access-list | policy> resources** command, refer to the “Displaying Current Resource Usage” section in the *Monitoring Resources* appendix.*

```
ProCurve# show policy resources
```

Includes the hardware resources used by classifier-based local and remote mirroring policies that are currently applied to interfaces on the switch.

Resource usage in Policy Enforcement Engine

Ports	Rules Available	Rules Used				VT	Mirror	Other
		ACL	QoS	IDM				
1-24	3014	15	11	0	1	0	3	
25-48	3005	15	10	10	1	0	3	
A	3017	15	8	0	1	0	3	

Ports	Meters Available	Meters Used				VT	Mirror	Other
		ACL	QoS	IDM				
1-24	250		5	0			0	
25-48	251		4	0			0	
A	253		2	0			0	

Ports	Application Port Ranges Available	Application Port Ranges Used				VT	Mirror	Other
		ACL	QoS	IDM				
1-24	3014	2	0	0			0	
25-48	3005	2	0	0			0	
A	3017	2	0	0			0	

0 of 8 Policy Engine management resources used.

Key:

- ACL = Access Control Lists
- QoS = Device & Application Port Priority, QoS Policies, ICMP rate limits
- IDM = Identity Driven Management
- VT = Virus Throttling blocks
- Mirror = Mirror Policies, Remote Intelligent Mirror endpoints
- Other = Management VLAN, DHCP Snooping, ARP Protection, Jumbo IP-MTU.

Resource usage includes resources actually in use, or reserved for future use by the listed feature. Internal dedicated-purpose resources, such as port bandwidth limits or VLAN QoS priority, are not included.

Figure B-45. Displaying the Hardware Resources Used by Currently Configured Mirroring Policies

Viewing the Mirroring Configurations in the Running Configuration File

Using the **show run** command, you can view the current mirroring configurations on the switch. In the **show run** command output, information about mirroring sources in configured sessions begins with the **mirror** keyword; monitored source interfaces are listed per-interface. For example:

```
ProCurve(config)# show run
Running configuration:
; J8697A Configuration Editor; Created on release #K.12.XX
max-vlans 300
ip access-list extended "100"
  10 permit icmp 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255 0
  exit
no ip address
exit
. . .
mirror 1 port B3
mirror 2 name "test-10" remote ip 10.10.10.1 8010 10.10.30.2
. . .
interface B1
  monitor ip access-group "100" In mirror 1
  monitor all Both mirror 2
  exit
. . .
```

Mirroring sessions with exit ports configured on the switch: B3 is an exit port for a local session; session 2 has a remote destination and exit port.

Selection criteria used to monitor traffic on port B1 for mirroring sessions 1 (ACL-based) and 2 (direction-based)

Figure B-46. Displaying Mirroring Sources and Sessions in the Running Configurations

Information about remote endpoints configured for remote sessions on the switch begin with the **mirror endpoint** keywords. In the following example, two remote sessions use the same exit port:

```
ProCurve(config)# show run
Running configuration:
; J8693A Configuration Editor; Created on release #K.12.XX
module 3 type J8694A
. . .
mirror endpoint ip 10.10.20.1 8010 10.10.30.2 port 4
mirror endpoint ip 10.10.51.10 7955 10.10.30.2 port 4
. . .
```

Remote endpoints configured on the switch, including source IP address, UDP port number, destination IP address, and remote exit port. Each remote session is identified by a unique UDP port number.

Figure B-47. Displaying Remote Mirroring Endpoints in the Running Configuration

Mirroring Configuration Examples

Example: Local Mirroring Using Traffic-Direction Criteria

An administrator wants to mirror the inbound traffic from workstation “X” on port A5 and workstation “Y” on port B17 to a traffic analyzer connected to port C24. In this case, the administrator chooses “1” as the session number. (Any unused session number from 1 to 4 is valid.) Since the switch provides both the source and destination for the traffic to monitor, local mirroring can be used. In this case, the command sequence is:

1. Configure the local mirroring session, including the exit port.
2. Configure the monitored source interfaces for the session.




Figure B-48. Local Mirroring Topology

```
ProCurve(config)# mirror 1 port c24
Caution: Please configure destination switch first.
Do you want to continue [y/n]? y
ProCurve(config)# interface a5,b17 monitor all in mirror 1
```

Configures port C24 as the mirroring destination (exit port) for session 1.

Reminder to configure mirroring destination before configuring source.

Mirrors all inbound and outbound traffic on ports A5 and B17 to the mirroring destination configured for session 1.

Figure B-49. Configuring a Local Mirroring Session for All Inbound and Outbound Port Traffic

Example: Remote Mirroring Using a Classifier-Based Policy

In the network shown in Figure B-50, an administrator has connected a traffic analyzer to port A15 (in VLAN 30) on switch C to monitor the TCP traffic to the server at 10.10.30.153 from workstations connected to switches A and B. Remote mirroring sessions are configured on switches A and B, and a remote mirroring endpoint on switch C. TCP traffic is routed through the network to the server from VLANs 10 and 20 on VLAN 30.

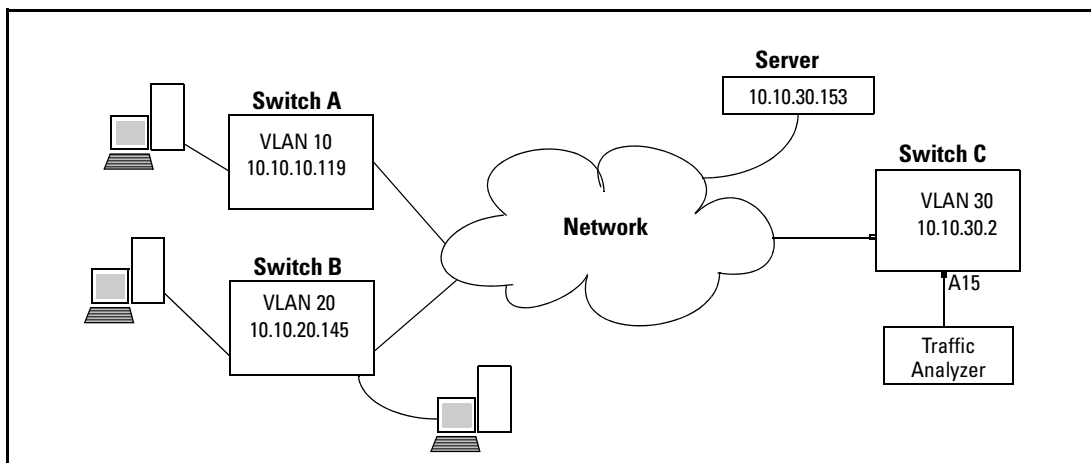


Figure B-50. Sample Topology in a Remote Mirroring Session

To configure this remote mirroring session using a classifier-based policy to select inbound TCP traffic on two VLAN interfaces, take the following steps:

1. On remote switch C, configure a remote mirroring endpoint using port A15 as the exit port (as described in “2. Configure a Mirroring Destination on a Remote Switch” on page B-50).




Figure B-51. Configuring a Remote Mirroring Endpoint: Remote Switch and Exit Port

2. On source switch A, configure an association between the remote mirroring endpoint on switch C and a mirroring session on switch A (as described in “3. Configure a Mirroring Session on the Source Switch” on page B-52).
3. On switch A, configure a classifier-based mirroring policy to select inbound TCP traffic destined to the server at 10.10.30.153, and apply the policy to the interfaces of VLAN 10 (as described in “Selecting Inbound Traffic Using Advanced Classifier-Based Mirroring” on page B-66).

On a source switch, associates session number 1 with a source IP address and UDP port, and a remote destination IP address.

```

1 Switch-A(config)# mirror [1] remote ip [10.10.10.119 9300] [10.10.30.2]
Caution: Please configure destination switch first.
Do you want to continue [y/n]? y

```

Class configuration that defines the matching TCP packets to be mirrored

```

2 Switch-A(config)# class ipv4 tcp7
Switch-A(class-config)# match tcp any 10.10.30.153
Switch-A(class-config)# match tcp any host 10.10.20.153/24
Switch-A(class-config)# match tcp any any eq 80
Switch-A(class-config)# exit

```

Policy configuration that defines the preconfigured class and session/destination device to which matching packets are mirrored

```

Switch-A(config)# policy mirror mirrorTCP
Switch-A(policy-config)# class ipv4 tcp7 action mirror 1
Switch-A(policy-config)# exit

```

Policy application to inbound traffic on a VLAN interface

```

3 Switch-A(config)# vlan 10 service-policy mirrorTCP in

```

1 The source IP address and UDP port number identify the mirroring source in session 1; the destination IP address identifies the remote switch to which traffic is mirrored. (The exit port for mirrored traffic, configured in Figure B-51, and the remote switch can belong to different VLANs.)

2 Configures a class that selects IPv4 TCP traffic destined to: the server at 10.10.30.153, a device in subnet 10.10.20.0, and any TCP traffic on port 80. (A packet that does not match these criteria is transmitted without being mirrored.)

3 Configures VLAN 10 as the source interface, and the mirroring policy as the selection criteria for inbound traffic on VLAN 10 in session 1.

Figure B-52. Configuring a Classifier-Based Policy on Source Switch A

4. On source switch B, repeat Steps 2 and 3:
 - a. Configure an association between the remote mirroring endpoint on switch C and a mirroring session on switch B.

- b. Configure a classifier-based mirroring policy to select inbound TCP traffic destined to the server at 10.10.30.153, and apply the policy to a VLAN interface for VLAN 20.

Because the remote session has mirroring sources on different switches, you can use the same session number (1) for both sessions.

The configuration of remote-mirroring session 1 on Switch B is the same as on Switch A (figure B-52), except for the difference in source VLAN and source IP address. Note that on different switches, the UDP port number (9300) can be the

```
Switch-B(config)# mirror 1 remote ip 10.10.20.145 9300 10.10.30.2
Caution: Please configure destination switch first.
Do you want to continue [y/n]? y
Switch-B(config)# class ipv4 tcp7
Switch-B(class-config)# match tcp any 10.10.30.153
Switch-B(class-config)# match tcp any host 10.10.20.153/24
Switch-B(class-config)# match tcp any any eq 80
Switch-B(class-config)# exit
Switch-B(config)# policy mirror mirrorTCP
Switch-B(policy-config)# class ipv4 tcp7 mirror 1
Switch-B(policy-config)# exit
Switch-B(config)# vlan 20 service-policy mirrorTCP in
```

Figure B-53. Configuring a Classifier-Based Policy on Source Switch B

Example: Remote Mirroring Using Traffic-Direction Criteria

In the network shown in Figure B-54, the administrator connects another traffic analyzer to port B10 (in VLAN 40) on switch C to monitor all traffic entering Switch A on port C12. For this mirroring configuration, the administrator configures a mirroring destination (with a remote exit port of B10) on switch C, and a remote mirroring session on Switch A.

If the mirroring configuration in the preceding example is enabled, it is necessary to use a different session number (2) and UDP port number (9400). (The IP address of the remote exit port [10.10.40.7] connected to traffic analyzer 2 [exit port B10] can belong to a different VLAN than the destination IP address of the VLAN used to reach remote switch C [10.20.40.1]).

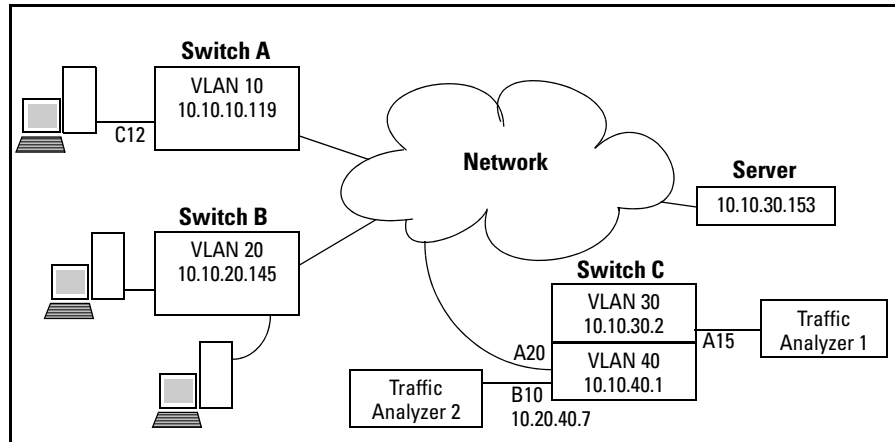


Figure B-54. Sample Topology for Remote Mirroring from a Port Interface

To configure this remote mirroring session using a directional-based traffic selection on a port interface, the operator must take the following steps:

1. On remote switch C, configure the remote mirroring endpoint using port B10 as the exit port for a traffic analyzer (as described in “2. Configure a Mirroring Destination on a Remote Switch” on page B-50):

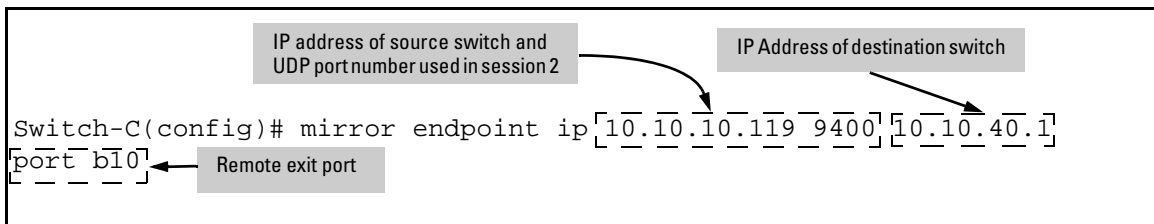


Figure B-55. Configuring a Remote Mirroring Endpoint

2. On source switch A, configure session 2 to use UDP port 9400 to reach the remote mirroring endpoint on switch C (10.10.40.1):
mirror 2 remote ip 10.10.10.119 9400 10.10.40.1
3. On source switch A, configure the local port C12 to select all inbound traffic to send to the preconfigured mirroring destination for session 2:
interface c12 monitor all in mirror 2

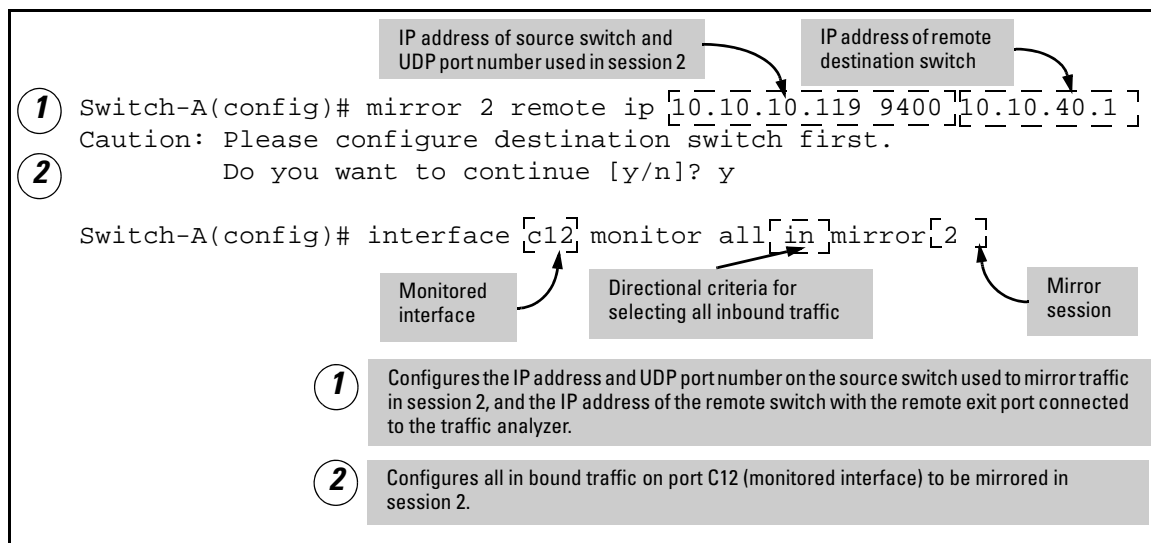


Figure B-56. Configuring a Remote Mirroring Session for Inbound Port Traffic

Maximum Supported Frame Size

The IPv4 encapsulation of mirrored traffic adds a 54-byte header to each mirrored frame. If a resulting frame exceeds the MTU (Maximum Transmission Unit) allowed in the network, the frame is dropped.

Note

Mirroring does not truncate frames, and oversized mirroring frames will be dropped. Also, remote mirroring does not allow downstream devices in a mirroring path to fragment mirrored frames.

If jumbo frames are enabled on the mirroring source switch, then the mirroring destination switch and all downstream devices connecting the source switch to the mirroring destination must be configured to support jumbo frames.

Enabling Jumbo Frames To Increase the Mirroring Path MTU

On 1 Gbps and 10 Gbps ports in the mirroring path, you can reduce the number of dropped frames by enabling jumbo frames on all intermediate switches and routers. (The maximum transmission unit—MTU—on the switches covered by this manual is 9220 bytes for frames having an 802.1Q VLAN tag, and 9216 bytes for untagged frames.) For information on configuring the switch for jumbo frames, refer to “Configuring Jumbo Frame Operation” on page 13-32.

Table B-2. Maximum Frame Sizes for Mirroring

	Frame Type Configuration	Maximum Frame Size	VLAN Tag	Frame Mirrored to Local Port	Frame Mirrored to Remote Port	
				Data	Data	IPv4 Header
Untagged	Non-Jumbo (default config.)	1518	0	1518	1464	54
	Jumbo ¹ on All VLANs	9216	0	9216	9162	54
	Jumbo ¹ On All But Source VLAN	1518	0	n/a ²	1464	54
Tagged	Non-Jumbo	1522	4	1522	1468	54
	Jumbo ¹ on All VLANs	9220	4	9218	9164	54
	Jumbo ¹ On All But Source VLAN	1522	4	n/a ²	1468	54

¹Jumbo frames are allowed on ports operating at or above 1 Gbps.
²For local mirroring, a non-Jumbo configuration on the source VLAN dictates an MTU of 1518 bytes for untagged frames, and an MTU of 1522 for tagged frames, regardless of the Jumbo configuration on any other VLANs on the switch.

Effect of Downstream VLAN Tagging on Untagged, Mirrored Traffic

In a remote mirroring application, if mirrored traffic leaves the switch without 802.1Q VLAN tagging, but is forwarded through a downstream device that adds 802.1Q VLAN tags, the MTU for untagged mirrored frames leaving the source switch is reduced below the values shown in Table B-2.

For example, if the MTU on the path to the destination is 1522 bytes, then untagged mirrored frames leaving the source switch cannot exceed 1518 bytes. Likewise, if the MTU on the path to the destination is 9220 bytes, then untagged mirrored frames leaving the source switch cannot exceed 9216 bytes.

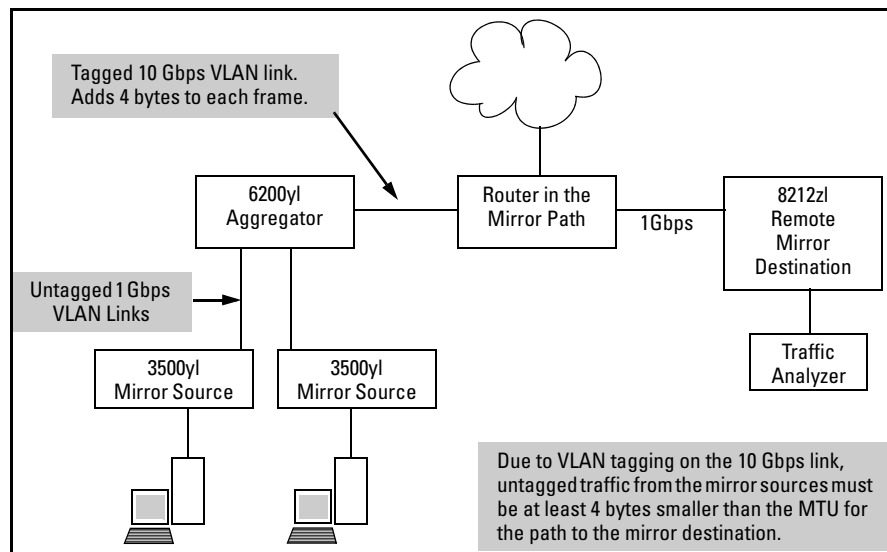


Figure B-57. Effect of Downstream VLAN Tagging on the MTU for Mirrored Traffic

Operating Notes for Traffic Mirroring

- **Mirroring Dropped Traffic:** When an interface is configured to mirror traffic to a local or remote destination, packets are mirrored regardless of whether the traffic is dropped while on the interface. For example, if an ACL is configured on a VLAN with a **deny** ACE that eliminates packets from a Telnet application, the switch still mirrors the Telnet packets that are received on the interface and subsequently dropped.
- **Mirroring and Spanning Tree:** Mirroring is performed regardless of the spanning-tree (STP) state of a port or trunk. This means, for example, that inbound traffic on a port blocked by STP can still be monitored for STP protocol packets during the STP setup phase.
- **Tagged and Untagged Frames:** For a frame entering or leaving the switch on a mirrored port, the mirrored copy retains the tagged or untagged state the original frame carried when it entered into or exited from the switch. (The tagged or untagged VLAN membership of ports in the path leading to the mirroring destination does not affect the tagged or untagged status of the mirrored copy itself.)

Thus, if a tagged frame arrives on a mirrored port, the mirrored copy will also be tagged, regardless of the status of ports in the destination path. If a frame exits from the switch on a mirrored port that is a tagged member of a VLAN, then the mirrored copy will also be tagged for the same reason.

To prevent a VLAN tag from being added to the mirrored copy of an outbound packet sent to a mirroring destination, you must enter the **no-tag-added** parameter when you configure a port, trunk, or mesh interface to select mirrored traffic. For more information see “Port Interface with Traffic Direction as the Selection Criteria” on page B-57 and “Untagged Mirrored Packets” on page B-59.

- **Effect of IGMP on Mirroring:** If both inbound and outbound mirroring is operating when IGMP is enabled on a VLAN, two copies of mirrored IGMP frames may appear at the mirroring destination.
- **Mirrored Traffic Not Encrypted:** Mirrored traffic undergoes IPv4 encapsulation, but mirrored encapsulated traffic is not encrypted.
- **IPv4 Header Added:** The IPv4 encapsulation of mirrored traffic adds a 54-byte header to each mirrored frame. If a resulting frame exceeds the maximum MTU allowed in the network, it will be dropped. To reduce the number of dropped frames, enable jumbo frames in the mirroring path, including all intermediate switches and/or routers. (The maximum transmission unit—MTU—on the switch is 9220 bytes, which includes 4 bytes for the 802.1Q VLAN tag.) For more information, refer to “Maximum Supported Frame Size” on page B-92. To configure the switch for jumbo frames, refer to “Configuring Jumbo Frame Operation” on page 13-32.

- **Intercepted or Injected Traffic:** The mirroring feature does not protect against either mirrored traffic being intercepted or traffic being injected into a mirrored stream by an intermediate host.
- **Inbound Mirrored IPv4-Encapsulated Frames are Not Mirrored:** The switch does not mirror IPv4-encapsulated mirrored frames that it receives on an interface. This prevents duplicate mirrored frames in configurations where the port connecting the switch to the network path for a mirroring destination is also a port whose inbound or outbound traffic is being mirrored. For example, if traffic leaving the switch through ports B5, B6, and B7 is being mirrored through port B7 to a network analyzer, the mirrored frames from traffic on ports B5 and B6 will not be mirrored a second time as they pass through port B7.
- **Switch Operation as Both Destination and Source:** A switch configured as remote destination switch can also be configured to mirror traffic to one of its own ports (local mirroring) or to a destination on another switch (remote mirroring).
- **Monitor Command Note:** If session 1 is already configured with a destination, you can enter the **[no] vlan < vid > monitor** or **[no] interface < port > monitor** command without mirroring criteria and a mirror session number. In this case, the switch automatically configures or removes mirroring for inbound and outbound traffic from the specified VLAN or port(s) to the destination configured for session 1.
- **Loss of Connectivity Suspends Remote Mirroring:** When a remote mirroring session is configured on a source switch, the switch sends an ARP request to the configured destination approximately every 60 seconds. If the source switch fails to receive the expected ARP response from the destination for the session, transmission of mirrored traffic in the session halts. However, because the source switch continues to send ARP requests for each configured remote session, link restoration or discovery of another path to the destination enables the source switch to resume transmitting the session's mirrored traffic after a successful ARP response cycle occurs. Note that if a link's connectivity is repeatedly interrupted ("link toggling"), little or no mirrored traffic may be allowed for sessions using that link. To verify the status of any mirroring session configured on the source switch, use the **show monitor** command.

Troubleshooting Traffic Mirroring

If mirrored traffic does not reach the configured remote destination (endpoint) switch or remote exit port, check the following configurations:

- In a remote mirroring session, the **mirror remote ip** command parameters configured on the source switch for source IP address, source UDP port, and destination IP address must be identical to the same parameters configured with the **mirror endpoint ip** command on the remote destination switch.
- The configured remote exit port must not be a member of a trunk or mesh.
- If the destination for mirrored traffic is on a different VLAN than the source, routing must be correctly configured along the path from the source to the destination.
- On the remote destination (endpoint) switch, the IP addresses of the remote exit port and the switch can belong to different VLANs.
- All links on the path from the source switch to the destination switch must be active.

Caution

A mirroring exit port should be connected only to a network analyzer, IDS, or other network edge device that has no connection to other network resources. Configuring a mirroring exit port connection to a network can result in serious network performance problems, and is strongly discouraged by ProCurve Networking.

Monitoring and Analyzing Switch Operation
Traffic Mirroring

Troubleshooting

Contents

Overview	C-4
Troubleshooting Approaches	C-5
Browser or Telnet Access Problems	C-6
Unusual Network Activity	C-8
General Problems	C-8
802.1Q Prioritization Problems	C-9
ACL Problems	C-9
IGMP-Related Problems	C-14
LACP-Related Problems	C-14
Mesh-Related Problems	C-15
Port-Based Access Control (802.1X)-Related Problems	C-15
QoS-Related Problems	C-18
Radius-Related Problems	C-18
Spanning-Tree Protocol (MSTP) and Fast-Uplink Problems	C-19
SSH-Related Problems	C-20
TACACS-Related Problems	C-22
TimeP, SNTP, or Gateway Problems	C-24
VLAN-Related Problems	C-24
Fan Failure	C-26
Using the Event Log for Troubleshooting Switch Problems	C-27
Event Log Entries	C-27
Menu: Displaying and Navigating in the Event Log	C-35
CLI: Displaying the Event Log	C-36
CLI: Clearing Event Log Entries	C-37
CLI: Turning Event Numbering On	C-37

Using Log Throttling to Reduce Duplicate Event Log and SNMP Messages	C-37
Log Throttle Periods	C-38
Example of Log Throttling	C-38
Example of Event Counter Operation	C-40
Debug/Syslog Operation	C-41
Debug/Syslog Messaging	C-41
Debug/Syslog Destination Devices	C-41
Debug/Syslog Configuration Commands	C-42
Configuring Debug/Syslog Operation	C-44
Displaying a Debug/Syslog Configuration	C-46
Debug Command	C-50
Debug Messages	C-50
Debug Destinations	C-52
Logging Command	C-54
Configuring a Syslog Server	C-55
Adding a Description for a Syslog Server	C-57
Adding a Priority Description	C-58
Configuring the Severity Level for Event Log Messages Sent to a Syslog Server	C-59
Configuring the System Module Used to Select the Event Log Messages Sent to a Syslog Server	C-60
Operating Notes for Debug and Syslog	C-60
Diagnostic Tools	C-62
Port Auto-Negotiation	C-63
Ping and Link Tests	C-63
Web: Executing Ping or Link Tests	C-64
CLI: Ping Test	C-65
Link Tests	C-66
Traceroute Command	C-67
Viewing Switch Configuration and Operation	C-71
CLI: Viewing the Startup or Running Configuration File	C-71
Web: Viewing the Configuration File	C-71
CLI: Viewing a Summary of Switch Operational Data	C-71
Saving show tech Command Output to a Text File	C-73

Customizing show tech Command Output	C-74
CLI: Viewing More Information on Switch Operation	C-78
Pattern Matching When Using the Show Command	C-79
CLI: Useful Commands for Troubleshooting Sessions	C-82
Restoring the Factory-Default Configuration	C-83
CLI: Resetting to the Factory-Default Configuration	C-83
Clear/Reset: Resetting to the Factory-Default Configuration	C-83
Restoring a Flash Image	C-84
DNS Resolver	C-87
Terminology	C-87
Basic Operation	C-88
Configuring and Using DNS Resolution with DNS-Compatible Commands	C-89
Configuring a DNS Entry	C-90
Example Using DNS Names with Ping and Traceroute	C-91
Viewing the Current DNS Configuration	C-93
Operating Notes	C-94
Event Log Messages	C-95
Locator LED (Locating a Switch)	C-96

Overview

This appendix addresses performance-related network problems that can be caused by topology, switch configuration, and the effects of other devices or their configurations on switch operation. (For switch-specific information on hardware problems indicated by LED behavior, cabling requirements, and other potential hardware-related problems, refer to the *Installation Guide* you received with the switch.)

Note

ProCurve periodically places switch software updates on the ProCurve Networking web site. ProCurve recommends that you check this web site for software updates that may have fixed a problem you are experiencing.

For information on support and warranty provisions, refer to the Support and Warranty booklet shipped with the switch.

Troubleshooting Approaches

Use these approaches to diagnose switch problems:

- Check the ProCurve Networking web site for software updates that may have solved your problem: **www.procurve.com**
- Check the switch LEDs for indications of proper switch operation:
 - Each switch port has a Link LED that should light whenever an active network device is connected to the port.
 - Problems with the switch hardware and software are indicated by flashing the Fault and other switch LEDs.

Refer to the *Installation Guide* shipped with the switch for a description of the LED behavior and information on using the LEDs for troubleshooting.

- Check the network topology/installation. Refer to the *Installation Guide* shipped with the switch for topology information.
- Check cables for damage, correct type, and proper connections. You should also use a cable tester to check your cables for compliance to the relevant IEEE 802.3 specification. Refer to the *Installation Guide* shipped with the switch for correct cable types and connector pin-outs.
- Use ProCurve Manager to help isolate problems and recommend solutions.
- Use the Port Utilization Graph and Alert Log in the web browser interface included in the switch to help isolate problems. Refer to Chapter 5, “Using the ProCurve Web Browser Interface” for operating information. These tools are available through the web browser interface:
 - Port Utilization Graph
 - Alert Log
 - Port Status and Port Counters screens
 - Diagnostic tools (Link test, Ping test, configuration file browser)
- For help in isolating problems, use the easy-to-access switch console built into the switch or Telnet to the switch console. Refer to chapters 3 and 4 for operating information on the Menu and CLI interfaces included in the console. These tools are available through the switch console
 - Status and Counters screens
 - Event Log
 - Diagnostics tools (Link test, Ping test, configuration file browser, and advanced user commands)

Browser or Telnet Access Problems

Cannot access the web browser interface:

- Access may be disabled by the **Web Agent Enabled** parameter in the switch console. Check the setting on this parameter by selecting:

2. Switch Configuration ...

1. System Information

- The switch may not have the correct IP address, subnet mask or gateway. Verify by connecting a console to the switch's Console port and selecting:

2. Switch Configuration ...

5. IP Configuration

Note: If DHCP/Bootp is used to configure the switch, the IP addressing can be verified by selecting:

1. Status and Counters ...

2. Switch Management Address Information

also check the DHCP/Bootp server configuration to verify correct IP addressing.

- If you are using DHCP to acquire the IP address for the switch, the IP address "lease time" may have expired so that the IP address has changed. For more information on how to "reserve" an IP address, refer to the documentation for the DHCP application that you are using.
- If one or more IP-Authorized managers are configured, the switch allows web browser access only to a device having an authorized IP address. For more information on IP Authorized managers, refer to the *Access Security Guide* for your switch.
- Java™ applets may not be running on the web browser. They are required for the switch web browser interface to operate correctly. Refer to the online Help on your web browser for how to run the Java applets.

Cannot Telnet into the switch console from a station on the network:

- Off subnet management stations can lose Telnet access if you enable routing without first configuring a static (default) route. That is, the switch uses the IP default gateway only while operating as a Layer 2 device. While routing is enabled on the switch, the IP default gateway is not used. You can avoid this problem by using the `ip route` command to configure a static (default) route before enabling routing. For more information, refer to the chapter titled “IP Routing Features” in the *Multicast and Routing Guide* for your switch.
- Telnet access may be disabled by the **Inbound Telnet Enabled** parameter in the System Information screen of the menu interface:

2. Switch Configuration

1. System Information

- The switch may not have the correct IP address, subnet mask, or gateway. Verify by connecting a console to the switch’s Console port and selecting:

2. Switch Configuration

5. IP Configuration

Note: If DHCP/Bootp is used to configure the switch, refer to the **Note**, above.

- If you are using DHCP to acquire the IP address for the switch, the IP address “lease time” may have expired so that the IP address has changed. For more information on how to “reserve” an IP address, refer to the documentation for the DHCP application that you are using.
- If one or more IP-Authorized managers are configured, the switch allows inbound telnet access only to a device having an authorized IP address. For more information on IP Authorized managers, refer to the *Access Security Guide* for your switch.

Unusual Network Activity

Network activity that fails to meet accepted norms may indicate a hardware problem with one or more of the network components, possibly including the switch. Such problems can also be caused by a network loop or simply too much traffic for the network as it is currently designed and implemented. Unusual network activity is usually indicated by the LEDs on the front of the switch or measured with the switch console interface or with a network management tool such as ProCurve Manager. Refer to the *Installation Guide* you received with the switch for information on using LEDs to identify unusual network activity.

A topology loop can also cause excessive network activity. The Event Log “FFI” messages can be indicative of this type of problem.

General Problems

The network runs slow; processes fail; users cannot access servers or other devices. Broadcast storms may be occurring in the network. These may be due to redundant links between nodes.

- If you are configuring a port trunk, finish configuring the ports in the trunk before connecting the related cables. Otherwise you may inadvertently create a number of redundant links (i.e. topology loops) that will cause broadcast storms.
- Turn on Spanning Tree Protocol to block redundant links (i.e. topology loops)
- Check for FFI messages in the Event Log.

Duplicate IP Addresses. This is indicated by this Event Log message:

ip: Invalid ARP source: IP address on IP address

where: both instances of *IP address* are the same address, indicating the switch’s IP address has been duplicated somewhere on the network.

Duplicate IP Addresses in a DHCP Network. If you use a DHCP server to assign IP addresses in your network and you find a device with a valid IP address that does not appear to communicate properly with the server or other devices, a duplicate IP address may have been issued by the server. This can occur if a client has not released a DHCP-assigned IP address after the intended expiration time and the server “leases” the address to another device.

This can also happen, for example, if the server is first configured to issue IP addresses with an unlimited duration, then is subsequently configured to issue IP addresses that will expire after a limited duration. One solution is to configure “reservations” in the DHCP server for specific IP addresses to be assigned to devices having specific MAC addresses. For more information, refer to the documentation for the DHCP server.

One indication of a duplicate IP address in a DHCP network is this Event Log message:

```
ip: Invalid ARP source: < IP-address > on <IP-address >
```

where: both instances of *IP-address* are the same address, indicating the IP address that has been duplicated somewhere on the network.

The Switch Has Been Configured for DHCP/Bootp Operation, But Has Not Received a DHCP or Bootp Reply. When the switch is first configured for DHCP/Bootp operation, or if it is rebooted with this configuration, it immediately begins sending request packets on the network. If the switch does not receive a reply to its DHCP/Bootp requests, it continues to periodically send request packets, but with decreasing frequency. Thus, if a DHCP or Bootp server is not available or accessible to the switch when DHCP/Bootp is first configured, the switch may not immediately receive the desired configuration. After verifying that the server has become accessible to the switch, reboot the switch to re-start the process.

802.1Q Prioritization Problems

Ports configured for non-default prioritization (level 1 - 7) are not performing the specified action. If the ports were placed in a trunk group after being configured for non-default prioritization, the priority setting was automatically reset to zero (the default). Ports in a trunk group operate only at the default priority setting.

ACL Problems

ACLs are properly configured and assigned to VLANs, but the switch is not using the ACLs to filter IP layer 3 packets.

1. The switch may be running with IP routing disabled. To ensure that IP routing is enabled, execute **show running** and look for the IP routing statement in the resulting listing. For example:

S

```
ProCurve(config)# show running
Running configuration:
; J8697A Configuration Editor; Created on release # K.11.00
hostname " HPswitch"
module 1 type J8702A
ip default-gateway 10.30.248.1
ip routing
logging 10.28.227.2
snmp-server community "public" Unrestricted
ip access-list extended "Controls for VLAN 20"
  permit tcp 0.0.0.0 255.255.255.255 10.10.20.98 0.0.0.0 eq 80
  permit tcp 0.0.0.0 255.255.255.255 10.10.20.21 0.0.0.0 eq 80
  deny tcp 0.0.0.0 255.255.255.255 10.10.20.1 0.0.0.255 eq 80
  deny tcp 10.10.20.17 0.0.0.0 10.10.10.100 0.0.0.0 eq 23 log
  deny tcp 10.10.20.23 0.0.0.0 10.10.10.100 0.0.0.0 eq 23 log
  deny tcp 10.10.20.40 0.0.0.0 10.10.10.100 0.0.0.0 eq 23 log
  permit ip 10.10.20.1 0.0.0.255 10.10.10.100 0.0.0.0
  deny ip 10.10.30.1 0.0.0.255 10.10.10.100 0.0.0.0
  permit ip 10.10.30.1 0.0.0.255 10.10.10.1 0.0.0.255
exit
-- MORE --, next page: Space, next line: Enter, quit: Control-C
```

Indicates that routing is enabled; a requirement for ACL operation. (There is an exception. Refer to the Note, below.)

Figure C-1. Indication that Routing Is Enabled

Note

If an ACL assigned to a VLAN includes an ACE referencing an IP address on the switch itself as a packet source or destination, the ACE screens traffic to or from this switch address regardless of whether IP routing is enabled. This is a security measure designed to help protect the switch from unauthorized management access.

If you need to configure IP routing, execute the **ip routing** command.

2. ACL filtering on the switches covered in this guide applies only to routed packets and packets having a destination IP address (DA) on the switch itself. Also, the switch applies assigned ACLs only at the point where traffic enters or leaves the switch on a VLAN. Ensure that you have correctly applied your ACLs (“in” and/or “out”) to the appropriate VLAN(s).

The switch does not allow management access from a device on the same VLAN.

The implicit **deny any** function that the switch automatically applies as the last entry in any ACL always blocks packets having the same DA as the switch’s IP address on the same VLAN. That is, bridged packets with the switch itself as the destination are blocked as a security measure. To preempt this action, edit the ACL to include an ACE that permits access to the switch’s DA on that VLAN from the management device.

Error (Invalid input) when entering an IP address.

When using the “host” option in the command syntax, ensure that you are not including a mask in either dotted decimal or CIDR format. Using the “host” option implies a specific host device and therefore does not permit any mask entry.

```

ProCurve(config)# access-list 6 permit host 10.28.100.100 ← Correct.
ProCurve(config)# access-list 6 permit host 10.28.100.100 [255.255.255.255]
Invalid input: 255.255.255.255
ProCurve(config)# access-list 6 permit host 10.28.100.100/32,
Invalid input: 10.28.100.100/32
  
```

Incorrect. No mask needed to specify a single host.

Figure C-2. Examples of Correctly and Incorrectly Specifying a Single Host

Apparent failure to log all “Deny” Matches.

Where the **log** statement is included in multiple ACEs configured with a “deny” option, a large volume of “deny” matches generating logging messages in a short period of time can impact switch performance. If it appears that the switch is not consistently logging all “deny” matches, try reducing the number of logging actions by removing the **log** statement from some ACEs configured with the “deny” action.

The switch does not allow any routed access from a specific host, group of hosts, or subnet.

The implicit **deny any** function that the switch automatically applies as the last entry in any ACL may be blocking all access by devices not specifically permitted by an entry in an ACL affecting those sources. If you are using the ACL to block specific hosts, a group of hosts, or a subnet, but want to allow any access not specifically permitted, insert **permit any** as the last explicit entry in the ACL.

The switch is not performing routing functions on a VLAN

Two possible causes of this problem are:

- Routing is not enabled. If **show running** indicates that routing is not enabled, use the **ip routing** command to enable routing.
- *On a switch covered in this guide*, an ACL may be blocking access to the VLAN. Ensure that the switch’s IP address on the VLAN is not blocked by one of the ACE entries in an ACL applied to that VLAN. A

common mistake is to either not explicitly permit the switch's IP address as a DA or to use a wildcard ACL mask in a deny statement that happens to include the switch's IP address. For an example of this problem, refer to the section titled "General ACL Operating Notes" in the "Access Control Lists (ACLs)" chapter of the latest *Access Security Guide* for your switch.

Routing Through a Gateway on the Switch Fails

Configuring a "deny" ACE that includes a gateway address can block traffic attempting to use the gateway as a next-hop.

Remote Gateway Case. For example, configuring ACL "101" (below) and applying it outbound on VLAN 1 in Figure C-4 includes the router gateway (10.0.8.1) needed by devices on other networks. This can prevent the switch from sending ARP and other routing messages to the gateway router to support traffic from authorized remote networks.

<p>In Figure C-4, this ACE denies access to the 10 Net's 10.0.8.1 router gateway needed by the 20 Net. (Subnet mask is 255.255.255.0.)</p>	<pre>ProCurve(config)# show access-list config ip access-list extended "101" deny ip 0.0.0.0 255.255.255.255 10.0.8.30 0.0.0.255 permit ip 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255 exit</pre>
--	--

Figure C-3. Example of ACE Blocking an Entire Subnet

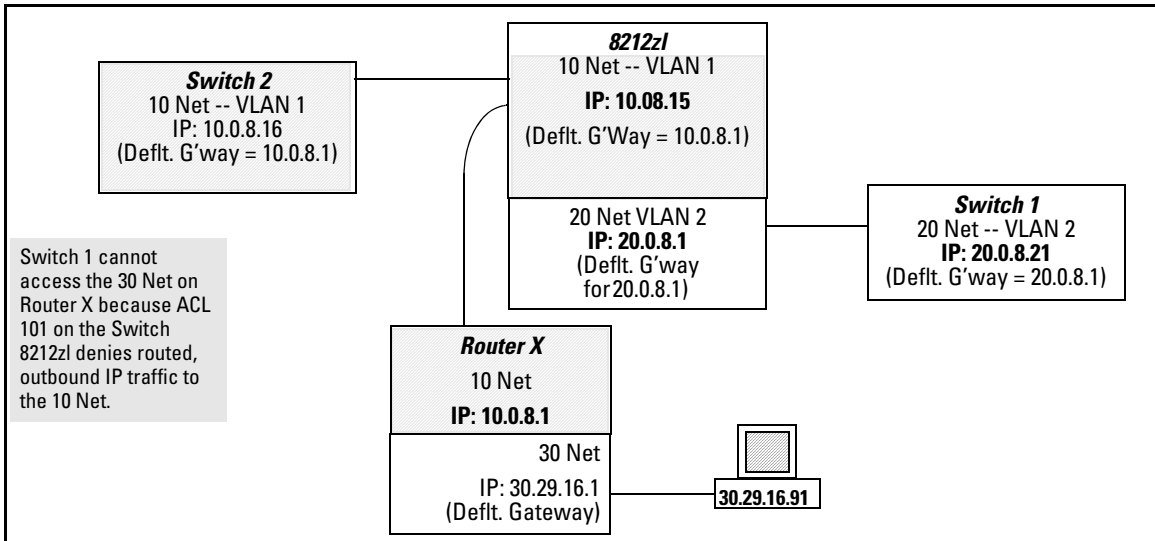


Figure C-4. Example of Inadvertently Blocking a Gateway

To avoid inadvertently blocking the remote gateway for authorized traffic from another network (such as the 20 Net in this example):

1. Configure an ACE that specifically permits authorized traffic from the remote network.
2. Configure narrowly defined ACEs to block unwanted IP traffic that would otherwise use the gateway. Such ACEs might deny traffic for a particular application, particular hosts, or an entire subnet.
3. Configure a “permit any” ACE to specifically allow any IP traffic to move through the gateway.

Local Gateway Case. If you use the switch as a gateway for traffic you want routed between subnets, use these general steps to avoid blocking the gateway for authorized applications:

1. Configure gateway security first for routing with specific permit and deny statements.
2. Permit authorized traffic.
3. Deny any unauthorized traffic that you have not already denied in step 1.

IGMP-Related Problems

IP Multicast (IGMP) Traffic That Is Directed By IGMP Does Not Reach IGMP Hosts or a Multicast Router Connected to a Port. IGMP must be enabled on the switch and the affected port must be configured for “Auto” or “Forward” operation.

IP Multicast Traffic Floods Out All Ports; IGMP Does Not Appear To Filter Traffic. The IGMP feature does not operate if the switch or VLAN does not have an IP address configured manually or obtained through DHCP/Bootp. To verify whether an IP address is configured for the switch or VLAN, do either of the following:

- **Try Using the Web Browser Interface:** If you can access the web browser interface, then an IP address is configured.
- **Try To Telnet to the Switch Console:** If you can Telnet to the switch, then an IP address is configured.
- **Using the Switch Console Interface:** From the Main Menu, check the Management Address Information screen by clicking on

1. Status and Counters

2. Switch Management Address Information

LACP-Related Problems

Unable to enable LACP on a port with the **interface < port-number > lacp** command. In this case, the switch displays the following message:

Operation is not allowed for a trunked port.

You cannot enable LACP on a port while it is configured as static **Trunk** port. To enable LACP on static-trunked port, first use the **no trunk < port-number >** command to disable the static trunk assignment, then execute **interface < port-number > lacp**.

Caution

Removing a port from a trunk without first disabling the port can create a traffic loop that can slow down or halt your network. Before removing a port from a trunk, ProCurve recommends that you either disable the port or disconnect it from the LAN.

Mesh-Related Problems

Traffic on a dynamic VLAN does not get through the switch mesh .

GVRP enables dynamic VLANs. Ensure that all switches in the mesh have GVRP enabled.

Port-Based Access Control (802.1X)-Related Problems

Note

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

See also “Radius-Related Problems” on page C-18.

The switch does not receive a response to RADIUS authentication requests. In this case, the switch will attempt authentication using the secondary method configured for the type of access you are using (console, Telnet, or SSH).

There can be several reasons for not receiving a response to an authentication request. Do the following:

- Use **ping** to ensure that the switch has access to the configured RADIUS servers.
- Verify that the switch is using the correct encryption key (RADIUS secret key) for each server.
- Verify that the switch has the correct IP address for each RADIUS server.
- Ensure that the **radius-server timeout** period is long enough for network conditions.

The switch does not authenticate a client even though the RADIUS server is properly configured and providing a response to the authentication request. If the RADIUS server configuration for authenticating the client includes a VLAN assignment, ensure that the VLAN exists as a static VLAN on the switch. Refer to “How 802.1X Authentication Affects VLAN Operation” in the *Access Security Guide* for your switch.

During RADIUS-authenticated client sessions, access to a VLAN on the port used for the client sessions is lost. If the affected VLAN is configured as untagged on the port, it may be temporarily blocked on that port during an 802.1X session. This is because the switch has temporarily assigned another

VLAN as untagged on the port to support the client access, as specified in the response from the RADIUS server. Refer to “How 802.1X Authentication Affects VLAN Operation” in the *Access Security Guide* for your switch.

The switch appears to be properly configured as a supplicant, but cannot gain access to the intended authenticator port on the switch to which it is connected. If `aaa authentication port-access` is configured for Local, ensure that you have entered the local *login* (operator-level) username and password of the authenticator switch into the **identity** and **secret** parameters of the supplicant configuration. If instead, you enter the enable (manager-level) username and password, access will be denied.

The supplicant statistics listing shows multiple ports with the same authenticator MAC address. The link to the authenticator may have been moved from one port to another without the supplicant statistics having been cleared from the first port. Refer to “Note on Supplicant Statistics” in the chapter on Port-Based and User-Based Access Control in the *Access Security Guide* for your switch.

The show port-access authenticator <port-list> command shows one or more ports remain open after they have been configured with control unauthorized. 802.1X is not active on the switch. After you execute `aaa port-access authenticator active`, all ports configured with **control unauthorized** should be listed as **Closed**.

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

ProCurve(config)# aaa port-access authenticator active

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

PortA9 shows an “Open” status even though Access Control is set to **Unauthorized** (Force Auth). This is because the port-access authenticator has not yet been activated.

Figure C-5. Authenticator Ports Remain “Open” Until Activated

RADIUS server fails to respond to a request for service, even though the server's IP address is correctly configured in the switch. Use **show radius** to verify that the encryption key (RADIUS secret key) the switch is using is correct for the server being contacted. If the switch has only a global key configured, then it either must match the server key or you must configure a server-specific key. If the switch already has a server-specific key assigned to the server's IP address, then it overrides the global key and must match the server key.

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

Server IP Addr      Auth  Acct
                  Port  Port  Encryption Key
-----
10.33.18.119      1812 1813  119-only-key
```

Figure C-6. Displaying Encryption Keys

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

The authorized MAC address on a port that is configured for both 802.1X and port security either changes or is re-acquired after execution of aaa port-access authenticator < port-list > initialize. If the port is force-authorized with **aaa port-access authenticator <port-list> control authorized** command and port security is enabled on the port, then executing **initialize** causes the port to clear the learned address and learn a new address from the first packet it receives after you execute **initialize**.

A trunked port configured for 802.1X is blocked. If you are using RADIUS authentication and the RADIUS server specifies a VLAN for the port, the switch allows authentication, but blocks the port. To eliminate this problem, either remove the port from the trunk or reconfigure the RADIUS server to avoid specifying a VLAN.

QoS-Related Problems

Loss of communication when using VLAN-tagged traffic. If you cannot communicate with a device in a tagged VLAN environment, ensure that the device either supports VLAN tagged traffic or is connected to a VLAN port that is configured as **Untagged**.

Radius-Related Problems

The switch does not receive a response to RADIUS authentication requests. In this case, the switch will attempt authentication using the secondary method configured for the type of access you are using (console, Telnet, or SSH).

There can be several reasons for not receiving a response to an authentication request. Do the following:

- Use **ping** to ensure that the switch has access to the configured RADIUS server.
- Verify that the switch is using the correct encryption key for the designated server.
- Verify that the switch has the correct IP address for the RADIUS server.
- Ensure that the **radius-server timeout** period is long enough for network conditions.
- Verify that the switch is using the same UDP port number as the server.

RADIUS server fails to respond to a request for service, even though the server's IP address is correctly configured in the switch. Use **show radius** to verify that the encryption key the switch is using is correct for the server being contacted. If the switch has only a global key configured, then it either must match the server key or you must configure a server-specific key. If the switch already has a server-specific key assigned to the server's IP address, then it overrides the global key and must match the server key.

```

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

```

Server IP Addr	Auth Port	Acct Port	Encryption Key
10.33.18.119	1812	1813	119-only-key

Global RADIUS Encryption Key

Unique RADIUS Encryption Key for the RADIUS server at 10.33.18.119

Figure C-7. Examples of Global and Unique Encryption Keys

Spanning-Tree Protocol (MSTP) and Fast-Uplink Problems

Caution

If you enable MSTP, it is recommended that you leave the remainder of the MSTP parameter settings at their default values until you have had an opportunity to evaluate MSTP performance in your network. Because incorrect MSTP settings can adversely affect network performance, you should avoid making changes without having a strong understanding of how MSTP operates. To learn the details of MSTP operation, refer to the IEEE 802.1s standard.

Broadcast Storms Appearing in the Network. This can occur when there are physical loops (redundant links) in the topology. Where this exists, you should enable MSTP on all bridging devices in the topology in order for the loop to be detected.

STP Blocks a Link in a VLAN Even Though There Are No Redundant Links in that VLAN. In 802.1Q-compliant switches MSTP blocks redundant physical links even if they are in separate VLANs. A solution is to use only one, multiple-VLAN (tagged) link between the devices. Also, if ports are available, you can improve the bandwidth in this situation by using a port trunk. Refer to “Spanning Tree Operation with VLANs” in the chapter titled “Static Virtual LANs (VLANs)” in the *Advanced Traffic Management Guide* for your switch.

Fast-Uplink Troubleshooting. Some of the problems that can result from incorrect usage of Fast-Uplink MSTP include temporary loops and generation of duplicate packets.

Problem sources can include:

- Fast-Uplink is configured on a switch that is the MSTP root device.
- Either the **Hello Time** or the **Max Age** setting (or both) is too long on one or more switches. Return the **Hello Time** and Max Age settings to their default values (2 seconds and 20 seconds, respectively, on a switch).
- A “downlink” port is connected to a switch that is further away (in hop count) from the root device than the switch port on which fast-uplink MSTP is configured.
- Two edge switches are directly linked to each other with a fast-uplink (Mode = **Uplink**) connection.
- Fast uplink is configured on both ends of a link.
- A switch serving as a backup MSTP root switch has ports configured for fast-uplink MSTP and has become the root device due to a failure in the original root device.

SSH-Related Problems

Switch access refused to a client. Even though you have placed the client’s public key in a text file and copied the file (using the **copy tftp pub-key-file** command) into the switch, the switch refuses to allow the client to have access. If the source SSH client is an SSHv2 application, the public key may be in the PEM format, which the switch (SSHv1) does not interpret. Check the SSH client application for a utility that can convert the PEM-formatted key into an ASCII-formatted key.

Executing IP SSH does not enable SSH on the switch. The switch does not have a host key. Verify by executing `show ip host-public-key`. If you see the message

```
ssh cannot be enabled until a host key is configured
(use 'crypto' command).
```

then you need to generate an SSH key pair for the switch. To do so, execute **crypto key generate**. (Refer to “2. Generating the Switch’s Public and Private Key Pair” in the SSH chapter of the *Access Security Guide* for your switch.)

Switch does not detect a client's public key that does appear in the switch's public key file (show ip client-public-key). The client's public key entry in the public key file may be preceded by another entry that does not terminate with a new line (CR). In this case, the switch interprets the next sequential key entry as simply a comment attached to the preceding key entry. Where a public key file has more than one entry, ensure that all entries terminate with a new line (CR). While this is optional for the last entry in the file, not adding a new line to the last entry creates an error potential if you either add another key to the file at a later time or change the order of the keys in the file.

An attempt to copy a client public-key file into the switch has failed and the switch lists one of the following messages.

```
Download failed: overlength key in key file.
```

```
Download failed: too many keys in key file.
```

```
Download failed: one or more keys is not a valid RSA  
public key.
```

The public key file you are trying to download has one of the following problems:

- A key in the file is too long. The maximum key length is 1024 characters, including spaces. This could also mean that two or more keys are merged together instead of being separated by a <CR><LF>.
- There are more than ten public keys in the key file.
- One or more keys in the file is corrupted or is not a valid rsa public key.

Client ceases to respond (“hangs”) during connection phase. The switch does not support data compression in an SSH session. Clients will often have compression turned on by default, but will disable it during the negotiation phase. A client which does not recognize the compression-request FAILURE response may fail when attempting to connect. Ensure that compression is turned off before attempting a connection to prevent this problem.

TACACS-Related Problems

Event Log. When troubleshooting TACACS+ operation, check the switch's Event Log for indications of problem areas.

All Users Are Locked Out of Access to the Switch. If the switch is functioning properly, but no username/password pairs result in console or Telnet access to the switch, the problem may be due to how the TACACS+ server and/or the switch are configured. Use one of the following methods to recover:

- Access the TACACS+ server application and adjust or remove the configuration parameters controlling access to the switch.
- If the above method does not work, try eliminating configuration changes in the switch that have not been saved to flash (boot-up configuration) by causing the switch to reboot from the boot-up configuration (which includes only the configuration changes made prior to the last **write memory** command.) If you did not use **write memory** to save the authentication configuration to flash, then pressing the Reset button or cycling the power reboots the switch with the boot-up configuration.
- Disconnect the switch from network access to any TACACS+ servers and then log in to the switch using either Telnet or direct console port access. Because the switch cannot access a TACACS+ server, it will default to local authentication. You can then use the switch's local Operator or Manager username/password pair to log on.
- As a last resort, use the Clear/Reset button combination to reset the switch to its factory default boot-up configuration. Taking this step means you will have to reconfigure the switch to return it to operation in your network.

No Communication Between the Switch and the TACACS+ Server Application. If the switch can access the server device (that is, it can **ping** the server), then a configuration error may be the problem. Some possibilities include:

- The server IP address configured with the switch's **tacacs-server host** command may not be correct. (Use the switch's **show tacacs-server** command to list the TACACS+ server IP address.)

- The encryption key configured in the server does not match the encryption key configured in the switch (by using the **tacacs-server key** command). Verify the key in the server and compare it to the key configured in the switch. (Use **show tacacs-server** to list the global key. Use **show config** or **show config running** to list any server-specific keys.)
- The accessible TACACS+ servers are not configured to provide service to the switch.

Access Is Denied Even Though the Username/Password Pair Is Correct. Some reasons for denial include the following parameters controlled by your TACACS+ server application:

- The account has expired.
- The access attempt is through a port that is not allowed for the account.
- The time quota for the account has been exhausted.
- The time credit for the account has expired.
- The access attempt is outside of the time frame allowed for the account.
- The allowed number of concurrent logins for the account has been exceeded

For more help, refer to the documentation provided with your TACACS+ server application.

Unknown Users Allowed to Login to the Switch. Your TACACS+ application may be configured to allow access to unknown users by assigning them the privileges included in a *default user* profile. Refer to the documentation provided with your TACACS+ server application.

System Allows Fewer Login Attempts than Specified in the Switch Configuration. Your TACACS+ server application may be configured to allow fewer login attempts than you have configured in the switch with the **aaa authentication num-attempts** command.

TimeP, SNTP, or Gateway Problems

The Switch Cannot Find the Time Server or the Configured Gateway .

TimeP, SNTP, and Gateway access are through the primary VLAN, which in the default configuration is the DEFAULT_VLAN. If the primary VLAN has been moved to another VLAN, it may be disabled or does not have ports assigned to it.

VLAN-Related Problems

Monitor Port. When using the monitor port in a multiple VLAN environment, the switch handles broadcast, multicast, and unicast traffic output from the monitor port as follows:

- If the monitor port is configured for tagged VLAN operation on the same VLAN as the traffic from monitored ports, the traffic output from the monitor port carries the same VLAN tag.
- If the monitor port is configured for untagged VLAN operation on the same VLAN as the traffic from the monitored ports, the traffic output from the monitor port is untagged.
- If the monitor port is not a member of the same VLAN as the traffic from the monitored ports, traffic from the monitored ports does not go out the monitor port.

None of the devices assigned to one or more VLANs on an 802.1Q-compliant switch are being recognized. If multiple VLANs are being used on ports connecting 802.1Q-compliant devices, inconsistent VLAN IDs may have been assigned to one or more VLANs. For a given VLAN, the same VLAN ID must be used on all connected 802.1Q-compliant devices.

Link Configured for Multiple VLANs Does Not Support Traffic for One or More VLANs. One or more VLANs may not be properly configured as “Tagged” or “Untagged”. A VLAN assigned to a port connecting two 802.1Q-compliant devices must be configured the same on both ports. For example, VLAN_1 and VLAN_2 use the same link between switch “X” and switch “Y”.

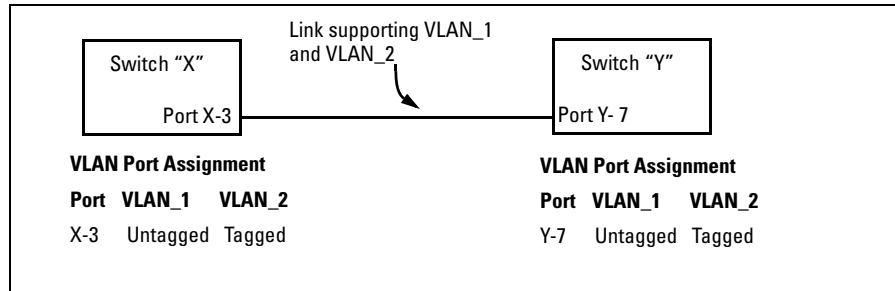


Figure C-8. Example of Correct VLAN Port Assignments on a Link

1. If VLAN_1 (VID=1) is configured as "Untagged" on port 3 on switch "X", then it must also be configured as "Untagged" on port 7 on switch "Y". Make sure that the VLAN ID (VID) is the same on both switches.
2. Similarly, if VLAN_2 (VID=2) is configured as "Tagged on the link port on switch "A", then it must also be configured as "Tagged" on the link port on switch "B". Make sure that the VLAN ID (VID) is the same on both switches.

Duplicate MAC Addresses Across VLANs. The switches covered in this guide operate with multiple forwarding databases. Thus, duplicate MAC addresses occurring on different VLANs can appear where a device having one MAC address is a member of more than one 802.1Q VLAN, and the switch port to which the device is linked is using VLANs (instead of MSTP or trunking) to establish redundant links to another switch. If the other device sends traffic over multiple VLANs, its MAC address will consistently appear in multiple VLANs on the switch port to which it is linked.

Note that attempting to create redundant paths through the use of VLANs will cause problems with some switches. One symptom is that a duplicate MAC address appears in the Port Address Table of one port, and then later appears on another port. While the switches have multiple forwarding databases, and thus does not have this problem, some switches with a single forwarding database for all VLANs may produce the impression that a connected device is moving among ports because packets with the same MAC address but different VLANs are received on different ports. You can avoid this problem by creating redundant paths using port trunks or spanning tree.

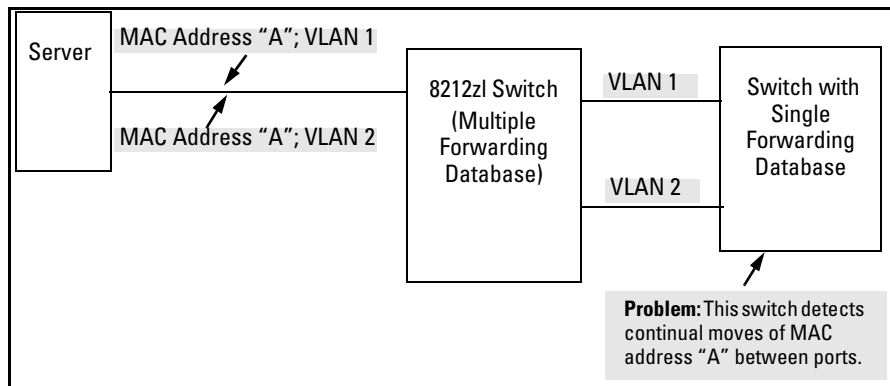


Figure C-9. Example of Duplicate MAC Address

Fan Failure

When two or more fans fail, a two-minute timer starts. After two minutes, the switch is powered down and must be rebooted to restart it. This protects the switch from possible overheating.

ProCurve recommends that you replace a failed fan tray assembly within one minute of removing it.

Using the Event Log for Troubleshooting Switch Problems

The Event Log records operating events in single- or double-line entries and serves as a tool to isolate and troubleshoot problems.

Starting in software release K.13.xx, the maximum number of entries supported in the Event Log is increased from 1000 to 2000 entries. Entries are listed in chronological order, from the oldest to the most recent.

Once the log has received 2000 entries, it discards the oldest message each time a new message is received. The Event Log window contains 14 log entry lines. You can scroll through it to view any part of the log.

Note

The Event Log is *erased* if power to the switch is interrupted or if you enter the **boot system** command. The contents of the Event Log are *not* erased if you:

- Reboot the switch by choosing the **Reboot Switch** option from the menu interface.
 - Enter the **reload** command from the CLI.
-

Event Log Entries

As shown in Figure C-10, each Event Log entry is composed of five or six fields, depending on whether numbering is turned on or not:

Severity	Date	Time	Event number	System Module	Event Message
I	08/05/06	10:52:32	00063	ports:	port A1 enabled

Figure C-10. Format of an Event Log Entry

Severity is one of the following codes (from highest to lowest severity):

- M** (major) indicates that a fatal switch error has occurred.
- E** (error) indicates that an error condition occurred on the switch.
- W** (warning) indicates that a switch service has behaved unexpectedly.

Troubleshooting

Using the Event Log for Troubleshooting Switch Problems

- I (information) provides information on normal switch operation.
- D (debug) is reserved for ProCurve internal diagnostic information.

Date is the date in the format *mm/dd/yy* when an entry is recorded in the log.

Time is the time in the format *hh:mm:ss* when an entry is recorded in the log.

Event Number is the number assigned to an event. You can turn event numbering on and off with the **[no] log-number** command.

System Module is the internal module (such as “ports:” for port manager) that generated a log entry. If VLANs are configured, then a VLAN name also appears for an event that is specific to an individual VLAN. Table C-1 lists the different system modules with a description of each one.

Event Message is a brief description of the operating event.

Table C-1. Event Log System Modules

System Module	Description	Documented in ProCurve Hardware/Software guide
802.1x	802.1X authentication: Provides access control on a per-client or per-port basis: <ul style="list-style-type: none">• Client-level security that allows LAN access to 802.1X clients (up to 32 per port) with valid user credentials• Port-level security that allows LAN access only on ports on which a single 802.1X-capable client (supplicant) has entered valid RADIUS user credentials	<i>Access Security Guide</i>
acl	Access Control Lists (ACLs): Filter layer-3 IP traffic to or from a host to block unwanted IP traffic, and block or limit other protocol traffic such as TCP, UDP, IGMP, and ICMP. Access control entries (ACEs) specify the filter criteria and an action (permit or deny) to take on a packet if it meets the criteria.	<i>Advanced Traffic Management Guide</i>
addrmgr	Address Table Manager: Manages MAC addresses that the switch has learned and are stored in the switch's address table.	<i>Management and Configuration Guide</i>
arp-protect	Dynamic ARP Protection: Protects the network from ARP cache poisoning. 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.	<i>Access Security Guide</i>
auth	Authorization: A connected client must receive authorization through web, AMC, RADIUS-based, TACACS+-based, or 802.1X authentication before it can send traffic to the switch.	<i>Access Security Guide</i>

System Module	Description	Documented in ProCurve Hardware/ Software guide
cdp	Cisco Discovery Protocol: Supports reading CDP packets received from neighbor devices, enabling a switch to learn about adjacent CDP devices. ProCurve switches do not support the transmission of CDP packets to neighbor devices.	<i>Management and Configuration Guide</i>
chassis	Hardware operation, including modules and ports, power supply, fans, transceivers, CPU interrupt errors, switch temperature, and so on. Chassis messages include events on Power Over Ethernet (POE) operation.	<i>Installation Guides</i> <i>Management and Configuration Guide</i>
connfilt	Connection-Rate filtering: Used on the network edge to protect the network from attack by worm-like malicious code by detecting hosts that are generating IP traffic that exhibits this behavior and (optionally) either throttling or dropping all IP traffic from the offending hosts. Connection-Rate filtering messages include events on virus throttling. Virus throttling uses connection-rate filtering to stop the propagation of malicious agents.	<i>Access Security Guide</i>
console	Console interface used to monitor switch and port status, reconfigure the switch, read the event log through an in-band Telnet or out-of-band connection.	<i>Installation and Getting Started Guide</i>
cos	Class of Service (CoS): Provides priority handling of packets traversing the switch, based on the IEEE 802.1p priority carried by each packet. CoS messages also include Quality of Service (QoS) events. The QoS feature classifies and prioritizes traffic throughout a network, establishing an end-to-end traffic priority policy to manage available bandwidth and improve throughput of important data.	<i>Advanced Traffic Management Guide</i>
dca	Dynamic Configuration Arbiter (DCA) determines the client-specific parameters that are assigned in an authentication session.	<i>Access Security Guide</i>
dhcp	Dynamic Host Configuration Protocol (DHCP) server configuration: Switch is automatically configured from a DHCP (Bootp) server, including IP address, subnet mask, default gateway, Timep Server address, and TFTP server address.	<i>Management and Configuration Guide</i>
dhcp v6c	DHCP for IPv6 prefix assignment	<i>IPv6 Configuration Guide</i>
dhcpr	DHCP relay: Forwards client-originated DHCP packets to a DHCP network server.	<i>Advanced Traffic Management Guide</i>
download	Download operation for copying a software version or files to the switch.	<i>Management and Configuration Guide</i>
dhcp-snoop	DHCP snooping: Protects your network from common DHCP attacks, such as address spoofing and repeated address requests.	<i>Access Security Guide</i>

Troubleshooting

Using the Event Log for Troubleshooting Switch Problems

System Module	Description	Documented in ProCurve Hardware/ Software guide
dma	Direct Access Memory (DMA): Transmits and receives packets between the CPU and the switch. Not used for logging messages in software release K.13.xx.	—
fault	Fault Detection facility, including response policy and the sensitivity level at which a network problem should generate an alert.	<i>Management and Configuration Guide</i>
ffi	Find, Fix, and Inform: Event or alert log messages indicating a possible topology loop that cause excessive network activity and results in the network running slow. FFI messages include events on transceiver connections with other network devices.	<i>Installation and Getting Started Guide</i> <i>Management and Configuration Guide</i>
garp	Generic Attribute Registration Protocol (GARP), defined in the IEEE 802.1D-1998 standard.	<i>Advanced Traffic Management Guide</i>
gvrp	GARP VLAN Registration Protocol (GVRP): Manages dynamic 802.1Q VLAN operations, in which the switch creates temporary VLAN membership on a port to provide a link to another port in the same VLAN on another device.	<i>Advanced Traffic Management Guide</i>
hpesp	Management module that maintains communication between switch ports.	<i>Installation and Getting Started Guide</i>
idm	Identity-driven Management: Optional management application used to monitor and control access to switch.	<i>Advanced Traffic Management Guide</i>
igmp	Internet Group Management Protocol: Reduces unnecessary bandwidth usage for multicast traffic transmitted from multimedia applications on a per-port basis.	<i>Multicast and Routing Guide</i>
inst-mon	Instrumentation Monitor: Identifies attacks on the switch by generating alerts for detected anomalies.	<i>Access Security Guide</i>
ip	IP addressing: Configures the switch with an IP address and subnet mask to communicate on the network and support remote management access; configures multiple IP addresses on a VLAN; enables IP routing on the switch.	<i>Management and Configuration Guide</i> <i>Multicast and Routing Guide</i>
ipaddrmgr	IP Address Manager: Programs IP routing information in switch hardware.	<i>Multicast and Routing Guide</i>
iplock	IP Lockdown: Prevents IP source address spoofing on a per-port and per-VLAN basis by forwarding only the IP packets in VLAN traffic that contain a known source IP address and MAC address binding for the port.	<i>Access Security Guide</i>
ipx	Novell Netware protocol filtering: On the basis of protocol type, the switch can forward or drop traffic to a specific set of destination ports on the switch.	<i>Access Security Guide</i>
licensing	ProCurve premium licensing: Provide access to expanded features on certain ProCurve network devices.	<i>Premium License Installation Guide</i>

System Module	Description	Documented in ProCurve Hardware/ Software guide
kms	Key Management System: Configures and maintains security information (keys) for all routing protocols, including a timing mechanism for activating and deactivating an individual protocol.	<i>Access Security Guide</i>
lACP	LACP trunks: The switch can either automatically establish an 802.3ad-compliant trunk group or provide a manually configured, static LACP trunk.	<i>Management and Configuration Guide</i>
ldbal	Load balancing in LACP port trunks or 802.1s Multiple Spanning Tree protocol (MSTP) that uses VLANs in a network to improve network resource utilization and maintain a loop-free environment. Load-balancing messages also include switch meshing events. The Switch Meshing feature provides redundant links, improved bandwidth use, and support for different port types and speeds.	<i>Management and Configuration Guide</i> <i>Advanced Traffic Management Guide</i>
lldp	Link-Layer Discovery Protocol: Supports transmitting LLDP packets to neighbor devices and reading LLDP packets received from neighbor devices, enabling a switch to advertise itself to adjacent devices and to learn about adjacent LLDP devices.	<i>Management and Configuration Guide</i>
loop_protect	Loop protection: Detects the formation of loops when an unmanaged device on the network drops spanning tree packets, and provides protection by transmitting loop protocol packets out ports on which loop protection has been enabled.	<i>Advanced Traffic Management Guide</i>
macauth	Web and MAC authentication: Port-based security employed on the network edge to protect private networks and the switch itself from unauthorized access using one of the following interfaces: <ul style="list-style-type: none"> • Web page login to authenticate users for access to the network • RADIUS server that uses a device's MAC address for authentication 	<i>Access Security Guide</i>
maclock	MAC lockdown and MAC lockout <ul style="list-style-type: none"> • MAC lockdown prevents station movement and MAC address "hijacking" by requiring a MAC address to be used only an assigned port on the switch. MAC Lockdown also restricts the client device to a specific VLAN. • MAC lockout blocks a specific MAC address so that the switch drops all traffic to or from the specified address. 	<i>Access Security Guide</i>
mgr	ProCurve Manager (PCM) and ProCurve Manager Plus (PCM+): Windows-based network management solutions for managing and monitoring performance of ProCurve devices. PCM messages also include events for configuration operations.	<i>Management and Configuration Guide</i>

Troubleshooting

Using the Event Log for Troubleshooting Switch Problems

System Module	Description	Documented in ProCurve Hardware/ Software guide
mld	Multicast Listener Discovery (MLD): IPv6 protocol used by a router to discover the presence of multicast listeners. MLD can also optimize IPv6 multicast traffic flow with the snooping feature.	<i>Multicast and Routing Guide</i>
mtm	Multicast Traffic Manager (MTM): Controls and coordinates L3 multicast traffic for upper layer protocols.	<i>Multicast and Routing Guide</i>
netinet	Network Internet: Monitors the creation of a route or an Address Resolution Protocol (ARP) entry and sends a log message in case of failure.	<i>Advanced Traffic Management Guide</i>
ospf	Open Short Path First (OSPF): A routing protocol that uses link-state advertisements (LSA) to update neighboring routers regarding its interfaces and information on those interfaces. Each routing switch maintains an identical database that describes its area topology to help a router determine the shortest path between it and any neighboring router.	<i>Multicast and Routing Guide</i>
pagp	Ports Aggregation Protocol (PAgP): Obsolete. Replaced by LACP (802.3ad). Not used for logging messages in software release K.13.xx.	—
pim	Protocol-independent multicast (PIM) routing: Enables IP multicast traffic to be transmitted for multimedia applications throughout a network without being blocked at routed interface (VLAN) boundaries.	<i>Multicast and Routing Guide</i>
ports	Port status and port configuration features, including mode (speed and duplex), flow control, broadcast limit, jumbo packets, and security settings. Port messages include events on Power Over Ethernet (POE) operation and transceiver connections with other network devices.	<i>Installation and Getting Started Guide</i> <i>Management and Configuration Guide</i> <i>Access Security Guide</i>
QinQ	IEEE 802.1ad specification, known as QinQ (provider bridging), provides a second tier of VLANs in a bridged network. QinQ supports the forwarding of traffic from multiple customers over a provider network using service VLANs (S-VLANs).	<i>Advanced Traffic Management Guide</i>
radius	RADIUS (Remote Authentication Dial-In User Service) authentication and accounting: A network server is used to authenticate user-connection requests on the switch and collect accounting information to track network resource usage.	<i>Access Security Guide</i>
ratelim	Rate-limiting: Enables a port to limit the amount of bandwidth a user or device may utilize for inbound traffic on the switch.	<i>Management and Configuration Guide</i>
sflow	Flow sampling: sFlow is an industry standard sampling technology, defined by RFC 3176, used to continuously monitor traffic flows on all ports providing network-wide visibility into the use of the network.	<i>Management and Configuration Guide</i>

System Module	Description	Documented in ProCurve Hardware/ Software guide
snmp	Simple Network Management Protocol: Allows you to manage the switch from a network management station, including support for security features, event reporting, flow sampling, and standard MIBs.	<i>Management and Configuration Guide</i>
sntp	Simple Network Time Protocol: Synchronizes and ensures a uniform time among interoperating devices.	<i>Management and Configuration Guide</i>
ssh	Secure Shell version 2 (SSHv2): Provides remote access to management functions on a switch via encrypted paths between the switch and management station clients capable of SSH operation. SSH messages also include events from the Secure File Transfer Protocol (SFTP) feature. SFTP provides a secure alternative to TFTP for transferring sensitive information, such as switch configuration files, to and from the switch in an SSH session.	<i>Access Security Guide</i>
ssl	Secure Socket Layer Version 3 (SSLv3), including Transport Layer Security (TLSv1) support: Provides remote web access to a switch via encrypted paths between the switch and management station clients capable of SSL/TLS operation.	<i>Access Security Guide</i>
stack	Stack management: Uses a single IP address and standard network cabling to manage a group (up to 16) of switches in the same IP subnet (broadcast domain), resulting in a reduced number of IP addresses and simplified management of small workgroups for scaling your network to handle increased bandwidth demand.	<i>Advanced Traffic Management Guide</i>
stp	Multiple-instance spanning tree protocol/MSTP (802.1s): Ensures that only one active path exists between any two nodes in a group of VLANs in the network. MSTP operation is designed to avoid loops and broadcast storms of duplicate messages that can bring down the network.	<i>Advanced Traffic Management Guide</i>
system	Switch management, including system configuration, switch bootup, activation of boot ROM image, memory buffers, traffic and security filters. System messages also include events from Management interfaces (menu, CLI, web browser, ProCurve Manager) used to reconfigure the switch and monitor switch status and performance.	<i>Management and Configuration Guide</i> <i>Access Security Guide</i>
tacacs	TACACS+ authentication: A central server is used to control access to the switches (and other TACACS-aware devices) in the network through a switch's console port (local access) or Telnet (remote access).	<i>Access Security Guide</i>
tcp	Transmission Control Protocol: A transport protocol that runs on IP and is used to set up connections.	<i>Advanced Traffic Management Guide</i>

Troubleshooting

Using the Event Log for Troubleshooting Switch Problems

System Module	Description	Documented in ProCurve Hardware/ Software guide
telnet	Session established on the switch from a remote device through the Telnet virtual terminal protocol.	<i>Management and Configuration Guide</i>
tftp	Trivial File Transfer Protocol: Supports the download of files to the switch from a TFTP network server.	<i>Management and Configuration Guide</i>
timep	Time Protocol: Synchronizes and ensures a uniform time among interoperating devices.	<i>Management and Configuration Guide</i>
udld	Uni-directional Link Detection: Monitors a link between two switches and blocks the ports on both ends of the link if the link fails at any point between the two devices.	<i>Access Security Guide</i>
udpf	UDP broadcast forwarding: Supports the forwarding of client requests sent as limited IP broadcasts addressed to a UDP application port on a network server.	<i>Multicast and Routing Guide</i>
update	Updates (TFTP or serial) to ProCurve software and updates to running-config and start-up config files	<i>Management and Configuration Guide</i>
usb	Auxiliary port that allows you to connect external devices to the switch.	<i>Installation and Getting Started Guide</i>
vlan	<p>Static 802.1Q VLAN operations, including port-and protocol-based configurations that group users by logical function instead of physical location</p> <ul style="list-style-type: none">• A port -based VLAN creates a layer-2 broadcast domain comprised of member ports that bridge IPv4 traffic among themselves.• A protocol-based VLAN creates a layer-3 broadcast domain for traffic of a particular routing protocol, and is comprised of member ports that bridge traffic of the specified protocol type among themselves. <p>VLAN messages include events from Management interfaces (menu, CLI, web browser, ProCurve Manager) used to reconfigure the switch and monitor switch status and performance.</p>	<i>Advanced Traffic Management Guide</i>
vrrp	Virtual Router Redundancy Protocol: Provides dynamic failover support as backup for gateway IP addresses (first-hop routers) so that if a VR's Master router becomes unavailable, the traffic it supports will be transferred to a backup router without major delays or operator intervention, eliminating single-point-of-failure problems.	<i>Advanced Traffic Management Guide</i>
xmodem	Xmodem: Binary transfer feature that supports the download of software files from a PC or Unix workstation.	<i>Management and Configuration Guide</i>
xrrp	Extended Router Redundancy Protocol: Routing protocol not used for logging messages in software release K.13.xx.	—

Menu: Displaying and Navigating in the Event Log

To display the Event Log from the Main Menu, select **Event Log**. Figure C-11 shows a sample event log display.

```
ProCurve Switch 5406z1                               25-Oct-2007  18:02:52
=====--CONSOLE - MANAGER MODE -----
M 10/25/07 16:30:02 sys: 'Operator cold reboot from CONSOLE session.'
I 10/25/07 17:42:51 00061 system: -----
I 10/25/07 17:42:51 00063 system: System went down:  10/25/07 16:30:02
I 10/25/07 17:42:51 00064 system: Operator cold reboot from CONSOLE session.
W 10/25/07 17:42:51 00374 chassis: WARNING: SSC is out of Date: Load 8.2 or newer
I 10/25/07 17:42:51 00068 chassis: Slot D Inserted
I 10/25/07 17:42:51 00068 chassis: Slot E Inserted
I 10/25/07 17:42:51 00068 chassis: Slot F Inserted
I 10/25/07 17:42:51 00690 udpf: DHCP relay agent feature enabled
I 10/25/07 17:42:51 00433 ssh: Ssh server enabled
I 10/25/07 17:42:52 00400 stack: Stack Protocol disabled
I 10/25/07 17:42:52 00128 tftp: Enable succeeded
I 10/25/07 17:42:52 00417 cdp: CDP enabled

----  Log events stored in memory 1-751.  Log events on screen 690-704.

Actions->   Back      Next page      Prev page      End      Help

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

Figure C-11. Example of an Event Log Display

The *log status line* below the recorded entries states the total number of events stored in the event log and which logged events are currently displayed.

To scroll to other entries in the Event Log, either preceding or following the currently visible portion, press the keys indicated at the bottom of the display (**Back**, **Next page**, **Prev page**, or **End**) or the keys described in Table C-1.

Table C-1. Event Log Control Keys

Key	Action
[N]	Advances the display by one page (next page).
[P]	Rolls back the display by one page (previous page).
[v]	Advances display by one event (down one line).

Key	Action
[^]	Rolls back display by one event (up one line).
[E]	Advances to the end of the log.
[H]	Displays Help for the Event Log.

CLI: Displaying the Event Log

To display messages recorded in the event log from the CLI, enter the **show logging** command. Keyword searches are supported.

Syntax: show logging [-a, -r] [<search-text>]

*By default, the **show logging** command displays the log messages recorded since the last reboot in chronological order.*

***-a** displays all recorded log messages, including those before the last reboot.*

***-r** displays all recorded log messages, with the most recent entries listed first.*

*<search-text> displays all Event Log entries that contain the specified text. Use a <search-text> value with **-a** or **-r** to further filter **show logging** command output.*

Examples. To display all Event Log messages that have “system” in the message text or module name, enter the following command:

```
ProCurve# show logging -a system
```

To display all Event Log messages recorded since the last reboot that have the word, “system”, in the message text or module name, enter:

```
ProCurve# show logging system
```

CLI: Clearing Event Log Entries

Use the **clear logging** command to hide, but not erase, Event Log entries displayed in **show logging** command output. Only new entries generated after you enter the command will be displayed.

To redisplay all hidden entries, including Event Log entries recorded prior to the last reboot, enter the **show logging -a** command.

Syntax: clear logging

Removes all entries from the event log display output.

CLI: Turning Event Numbering On

Syntax: [no] log-numbers

Turns event numbering on and off

Using Log Throttling to Reduce Duplicate Event Log and SNMP Messages

A recurring event can generate a series of duplicate Event Log messages and SNMP traps in a relatively short time. As a result, the Event Log and any configured SNMP trap receivers may be flooded with excessive, exactly identical messages. To help reduce this problem, the switch uses *log throttle periods* to regulate (throttle) duplicate messages for recurring events, and maintains a counter to record how many times it detects duplicates of a particular event since the last system reboot.

When the first instance of a particular event or condition generates a message, the switch initiates a log throttle period that applies to all recurrences of that event. If the logged event recurs during the log throttle period, the switch increments the counter initiated by the first instance of the event, but does not generate a new message.

If the logged event repeats again after the log throttle period expires, the switch generates a duplicate of the first message, increments the counter, and starts a new log throttle period during which any additional instances of the event are counted, but not logged. Thus, for a particular recurring event, the switch displays only one message in the Event Log for each log throttle period in which the event reoccurs. Also, each logged instance of the event message includes counter data showing how many times the event has occurred since the last reboot. The switch manages messages to SNMP trap receivers in the same way.

Log Throttle Periods

The length of the log throttle period differs according to an event's severity level:

Severity Level	Log Throttle Period
I (Information)	6000 Seconds
W (Warning)	600 Seconds
D (Debug)	60 Seconds
M (Major)	6 Seconds

Example of Log Throttling

For example, suppose that you configure VLAN 100 on the switch to support PIM operation, but do not configure an IP address. If PIM attempted to use VLAN 100, the switch would generate the first instance of the following Event Log message and counter.

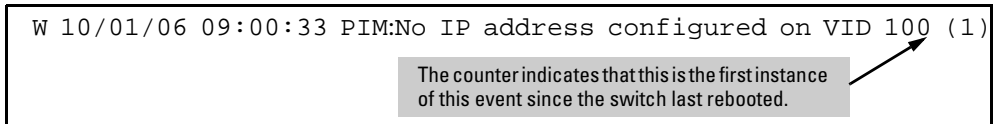


Figure C-12. Example of the First Instance of an Event Message and Counter

If PIM operation caused the same event to occur six more times during the initial log throttle period, there would be no further entries in the Event Log. However, if the event occurred again after the log throttle period expired, the switch would repeat the message (with an updated counter) and start a new log throttle period.

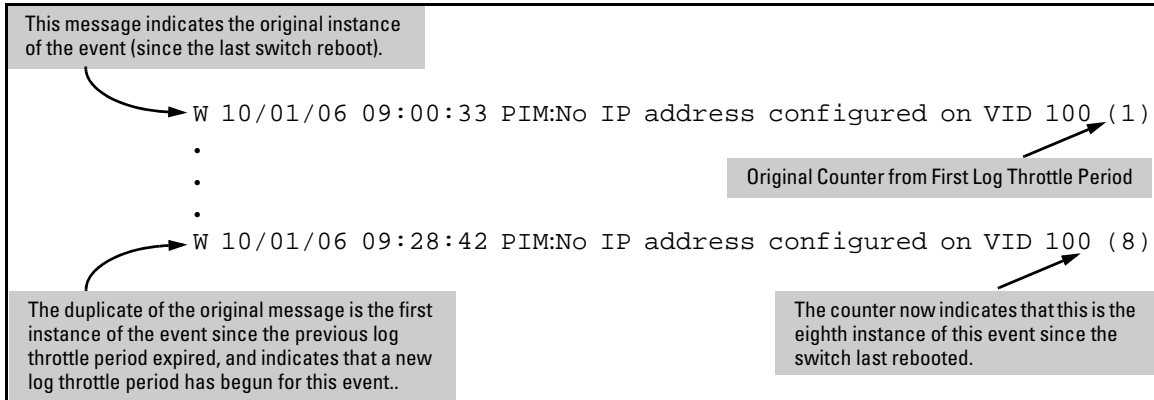


Figure C-13. Example of Duplicate Messages Over Multiple Log Throttling Periods

Note that if the same type of event occurs under different circumstances, the switch handles these as unrelated events for the purpose of Event Log messages. For example, if PIM operation simultaneously detected that VLANs 100 and 205 were configured without IP addresses, you would see log messages similar to the following:

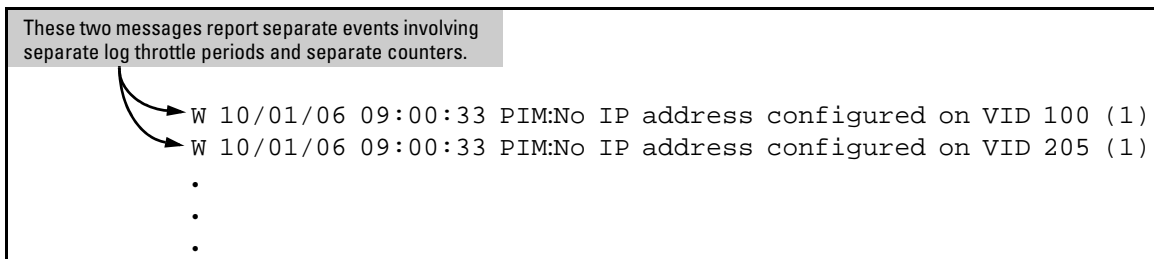


Figure C-14. Example of Log Messages Generated by Unrelated Events of the Same Type

Example of Event Counter Operation

Suppose the switch detects the following after a reboot:

- Three duplicate instances of the PIM “Send error” during the first log throttle period for this event
- Five more instances of the same Send error during the second log throttle period for this event
- Four instances of the same Send error during the third log throttle period for this event

In this case, the duplicate message would appear three times in the Event Log (once for each log throttle period for the event being described), and the Duplicate Message Counter would increment as shown in table C-2. (The same operation would apply for messages sent to any configured SNMP trap receivers.)

Table C-2. How the Duplicate Message Counter Increments

Instances During 1st Log Throttle Period	Instances During 2nd Log Throttle Period	Instances During 3rd Log Throttle Period	Duplicate Message Counter*
3			1
	5		4
		4	9

*This value always comprises the first instance of the duplicate message in the current log throttle period plus all previous occurrences of the duplicate message occurring since the switch last rebooted.

Debug/Syslog Operation

While the Event Log records switch-level progress, status, and warning messages on the switch, the Debug/System Logging (*Syslog*) feature provides a way to record Event Log and debug messages on a remote device. For example, you can send messages about routing misconfigurations and other network protocol details to an external device, and later use them to debug network-level problems.

Debug/Syslog Messaging

The Debug/Syslog feature allows you to specify the types of Event Log and debug messages that you want to send to an external device. As shown in Figure C-15, you can perform the following operations:

- Use the **debug** command to configure messaging reports for the following event types:
 - ACL “deny” matches
 - Dynamic ARP protection events
 - DHCP snooping events
 - Events recorded in the switch’s Event Log
 - IPv4 routing events
 - LLDP events
 - OSPF events
 - SSH events
 - VRRP events
 - Wireless Services events
- Use the **logging** command to select a subset of Event Log messages to send to an external device for debugging purposes according to:
 - Severity level
 - System module

Debug/Syslog Destination Devices

To use Debug/Syslog messaging, you must configure an external device as the logging destination by using the **logging** and **debug destination** commands. For more information, see “Debug Destinations” on page C-52 and “Configuring a Syslog Server” on page C-55.

A Debug/Syslog destination device can be a Syslog server and/or a console session. You can configure debug and logging messages to be sent to:

- Up to six Syslog servers
- A CLI session through a direct RS-232 console connection, or a Telnet or SSH session

Debug/Syslog Configuration Commands

Event Notification Logging	—	Automatically sends switch-level event messages to the switch's Event Log. Debug and Syslog do not affect this operation, but add the capability of directing Event Log messaging to an external device.
logging Command	<syslog-ip-addr>	Enables Syslog messaging to be sent to the specified IP address.
	facility	(Optional) The logging facility command specifies the destination (facility) subsystem used on a Syslog server for debug reports.
	severity	Sends Event Log messages of equal or greater severity than the specified value to configured debug destinations. (The default setting is to send Event Log messages from all severity levels.)
	system-module	Sends Event Log messages from the specified system module to configured debug destinations. The severity filter is also applied to the system-module messages you select. The default setting is to send Event Log messages from all system modules. To restore the default setting, enter the no logging system-module <system-module> or logging system-module all-pass commands.
debug Command	acl	Sends ACL Syslog logging to configured debug destinations. When there is a match with a "deny" statement, directs the resulting message to the configured debug destination(s).
	all	Sends debug logging to configured debug destinations for all ACL, Event Log, IP-OSPF, and IP-RIP options.
	arp-protect	Monitor and troubleshoot the validation of ARP packets
	destination	logging: Disables or re-enables Syslog logging on one or more Syslog servers configured with the logging < syslog-ip-addr > command. See "Debug Destinations" on page C-52. session: Assigns or re-assigns destination status to the terminal device that was most recently used to request debug output. "Debug Destinations" on page C-52. buffer: Enables Syslog logging to send the debug message types specified by the debug < debug-type > command to a buffer in switch memory. See "Debug Destinations" on page C-52. windshell: print debug messages to windshell.
	dhcp-snooping	agent: Displays DHCP snooping agent messages. event: Displays DHCP snooping event messages. packet: Displays DHCP snooping packet messages.
	dynamic-ip-lockdown	Displays dynamic IP lockdown messages.

event	Sends standard Event Log messages to configured debug destinations. (The same messages are also sent to the switch's Event Log, regardless of whether you enable this option.)
ip	forwarding: Sends IPv4 forwarding messages to the debug destination(s). ospf: Sends OSPF event logging to the debug destination(s). packet: Sends IPv4 packet messages to the debug destination(s). rip: Sends RIP event logging to the debug destination(s).
ipv6	dhcpv6-client: Sends DHCPv6 client debug messages to the configured debug destination. forwarding: Sends IPv6 forwarding messages to the debug destination(s) nd: Sends IPv6 debug messages for IPv6 neighbor discovery to the configured debug destination(s). packet: Sends IPv6 packet messages to the debug destination(s).
lldp	Sends LLDP debug logging to the debug destination(s).
ssh	Sends SSH debug messages at the specified level to the debug destination. The levels are fatal, error, info, verbose, debug, debug2, and debug3.
vrrp	Turns on tracing of the incoming and outgoing VRRP packets and sends debug logging to the debug destination.
wireless-services	Sends wireless service module debug messages to the debug destination.

Figure C-15. Summary of Debug/Syslog Configuration Commands

Using the Debug/Syslog feature, you can perform the following operations:

- Configure the switch to send Event Log messages to one or more Syslog servers. In addition, you can configure the messages to be sent to the User log facility (default) or to another log facility on configured Syslog servers.

Note

As of November 2008, the **logging facility** < *facility-name* > option (described on page C-57) is supported on the following switch models:

- 8200zl switches
- Series 6400cl switches
- 6200yl Switch
- 6600 switch
- Series 5400zl switches
- Series 5300xl switches
- Series 4200vl switches
- Series 4100gl switches (software release G.07.50 or greater)
- Series 3500 switches

- Series 3500yl switches
- Series 3400cl switches
- Series 2900 switches
- Series 2800 switches
- Series 2610 switches
- Series 2600 switches and the Switch 6108 (software release H.07.30 or greater)

For the latest feature information on ProCurve switches, visit the ProCurve Networking web site and check the latest release notes for the switch products you use.

-
- Configure the switch to send Event Log messages to the current management-access session (serial-connect CLI, Telnet CLI, or SSH).
 - Disable all Syslog debug logging while retaining the Syslog addresses from the switch configuration. This allows you to configure Syslog messaging and then disable and re-enable it as needed.
 - Display the current debug configuration. If Syslog logging is currently active, the list of configured Syslog servers is displayed.
 - Display the current Syslog server list when Syslog logging is disabled.

Configuring Debug/Syslog Operation

1. To use a Syslog server as the destination device for debug messaging, follow these steps:
 - a. Enter the **logging** *< syslog-ip-addr >* command at the global configuration level to configure the Syslog server IP address and enable Syslog logging. Optionally, you may also specify the destination subsystem to be used on the Syslog server by entering the **logging facility** command.

If no other Syslog server IP addresses are configured, entering the **logging** command enables both debug messaging to a Syslog server and the Event debug message type. As a result, the switch automatically sends Event Log messages to the Syslog server, regardless of other debug types that may be configured.
 - b. Re-enter the **logging** command in Step “a” to configure additional Syslog servers. You can configure up to a total of six servers. (When multiple server IP addresses are configured, the switch sends the debug message types that you configure in Step 3 to all IP addresses.)

2. To use a CLI session on a destination device for debug messaging:
 - a. Set up a serial, Telnet, or SSH connection to access the switch's CLI.
 - b. Enter the **debug destination session** command at the manager level.
3. Enable the types of debug messages to be sent to configured Syslog servers and/or the current session device by entering the **debug < debug-type >** command:

```
ProCurve# debug <acl|all|arp-protect|event|ip  
[bgp|forwarding|ospf|packet|rip|routemap]|ipv6|  
lldp|vrrp>
```

Repeat this step if necessary to enable multiple debug message types.

By default, Event Log messages are sent to configured debug destination devices. To block Event Log messages from being sent, enter the **no debug event** command.

4. If necessary, enable a subset of Event Log messages to be sent to configured Syslog servers by specifying a severity level and/or system module using the following commands

```
ProCurve(config)# logging severity <debug|major|error|warning|info>  
ProCurve(config)# logging system-module <system-module>
```

To display a list of valid values for each command, enter **logging severity** or **logging system-module** followed by **?** or pressing the **Tab** key.

The severity levels in order from the highest to lowest severity are: major, error, warning, info, debug. For a list of valid values for the **logging system-module <system-module >** command, refer to Table C-1 on page C-28.

5. If you configure system-module and/or severity-level values to filter Event Log messages, when you finish troubleshooting, you may want to reset these values to their default settings so that the switch sends all Event Log messages to configured debug destinations (Syslog servers and/or CLI session).

To remove a configured setting and restore the default values that send all Event Log messages, enter one or both of the following commands:

```
ProCurve(config)# no logging severity <debug|major|error|warning|info>  
ProCurve(config)# no logging system-module <system-module >
```

Caution

If you configure a severity-level, system-module, logging destination, or logging facility value and save the settings to the startup configuration (for example, by entering the **write memory** command), the debug settings are

saved after a system reboot (power cycle or reboot) and re-activated on the switch. As a result, after switch startup, one of the following situations may occur:

- Only a partial set of Event Log messages may be sent to configured debug destinations.
- Messages may be sent to a previously configured Syslog server used in an earlier debugging session.

Displaying a Debug/Syslog Configuration

Use the **show debug** command to display the currently configured settings for:

- Debug message types and Event Log message filters (severity level and system module) sent to debug destinations
- Debug destinations (Syslog servers or CLI session) and Syslog server facility to be used

Syntax: show debug

*Displays the currently configured debug logging destinations and message types selected for debugging purposes. (If no Syslog server address is configured with the **logging <syslog-ip-addr>** command, no **show debug** command output is displayed.)*

```
ProCurve(config)# show debug

Debug Logging
Destination:
Logging --
 10.28.38.164
Facility=kern
Severity=warning
System module=all-pass
Enabled debug types:
 event
```

Figure C-16. Sample Output of show debug Command

Example: In the following example, no Syslog servers are configured on the switch (default setting). When you configure a Syslog server, debug logging is enabled to send Event Log messages to the server. To limit the Event Log

messages sent to the Syslog server, specify a set of messages by entering the **logging severity** and **logging system-module** commands.

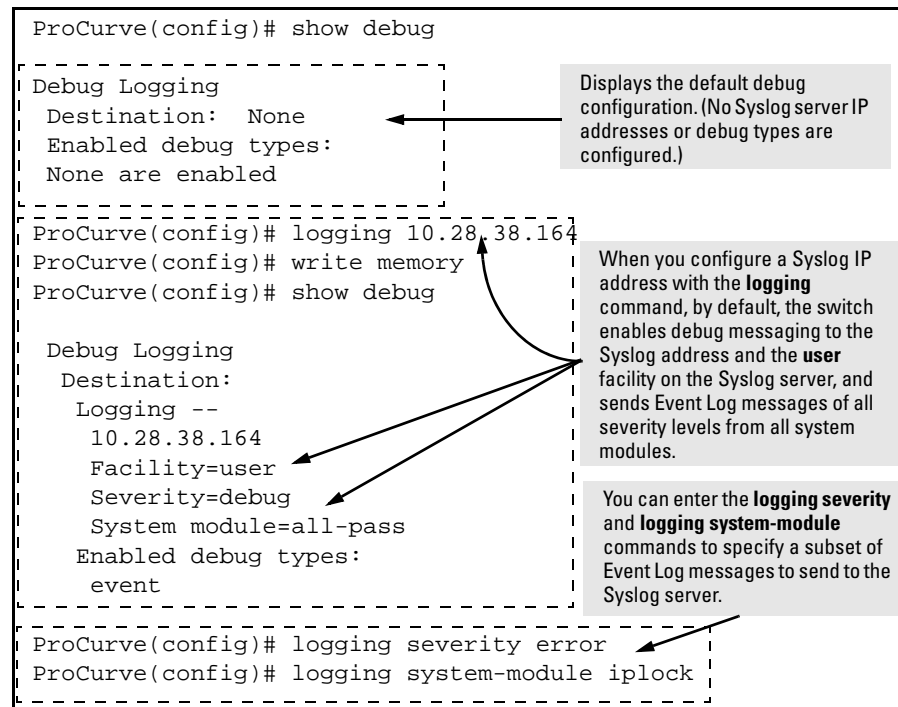


Figure C-17. Syslog Configuration to Receive Event Log Messages From Specified System Module and Severity Levels

As shown at the top of Figure C-17, if you enter the **show debug** command when no Syslog server IP address is configured, the configuration settings for Syslog server facility, Event Log severity level and system module are not displayed.

However, after you configure a Syslog server address and enable Syslog logging, all debug and logging settings are displayed with the **show debug** command. If you do not want Event Log messages sent to Syslog servers, you can block the messages from being sent by entering the **no debug event** command. (There is no effect on the normal logging of messages in the switch's Event Log.)

Example. The next example shows how to configure:

- Debug logging of ACL and IP-OSPF packet messages on a Syslog server at 18.38.64.164 (with **user** as the default logging facility).
- Display of these messages in the CLI session of your terminal device's management access to the switch.
- Blocking Event Log messages from being sent from the switch to the Syslog server and a CLI session.

To configure Syslog operation in these ways with the Debug/Syslog feature disabled on the switch, you would enter the commands shown in Figure C-18.

```
ProCurve# config
ProCurve(config)# logging 10.38.64.164
ProCurve(config)# show debug
  Debug Logging
  Destination:
  Logging --
    10.38.64.164
    Facility=user
    Severity=debug
    System module=all-pass
  Enabled debug types:
    event
ProCurve(config)# no debug event
ProCurve(config)# debug acl
ProCurve(config)# debug ip ospf packet
ProCurve(config)# debug destination session
ProCurve(config)# show debug
  Debug Logging
  Destination:
  Logging --
    10.38.64.164
    Facility=user
    Severity=debug
    System module=all-pass
  Session
  Enabled debug types:
    acl log
    ip ospf packet
```

Configure a Syslog server IP address. (No other Syslog servers are configured on the switch.) The server address serves as an active debug destination for any configured debug types.)

Display the new debug configuration. (Default debug settings - facility, severity, system module, and debug types- are displayed.)

Remove the unwanted event message logging to debug destinations.

Configure the debug messages types that you want to send to the Syslog server and CLI session.

Configure the CLI session as a debug destination.

Display the final debug and Syslog server configuration.

Figure C-18. Debug/Syslog Configuration for Multiple Debug Types and Multiple Destinations

Debug Command

At the manager level, use the **debug** command to perform two main functions:

- Specifies the types of event messages to be sent to an external destination.
- Specifies the destinations to which selected message types are sent.

By default, no debug destination is enabled and only Event Log messages are enabled to be sent.

Note

To configure a Syslog server, use the **logging** *<syslog-ip-addr>* command. For more information, see “Configuring a Syslog Server” on page C-55.

Debug Messages

Use the **debug** command to configure the types of debug messages that the switch can send to configured debug destinations.

Syntax: [no] debug *< debug-type >*

acl

*When a match occurs on an ACL “deny” Access Control Entry (with **log** configured), the switch sends an ACL message to configured debug destinations. For information on ACLs, refer to the “Access Control Lists (ACLs)” chapter in the latest version of the following guides:*

- *IPv4 ACLs: [Access Security Guide](#)*
- *IPv6 ACLs: [IPv6 Configuration Guide](#)*

Note: *Beginning with software release K.14.01, ACE matches (hits) for permit and deny entries can be tracked using the **show statistics < aclv4 | aclv6 >** command. (Default: Disabled - ACL messages for traffic that matches “deny” entries are not sent.)*

all

Configures the switch to send all debug message types to configured debug destination(s). (Default: Disabled - No debug messages are sent.)

event

Configures the switch to send Event Log messages to configured debug destinations.

Note: *This value does not affect the reception of event notification messages in the Event Log on the switch.*

Syntax: [no] debug < debug-type > (Continued)

event

Event Log messages are automatically enabled to be sent to debug destinations in these conditions:

- *If no Syslog server address is configured and you enter the **logging** <syslog-ip-addr> command to configure a destination address.*
- *If at least one Syslog server address is configured in the startup configuration and the switch is rebooted or reset.*

Event log messages are the default type of debug message sent to configured debug destinations.

ip

Enables all IP-OSPF messages for configured destinations.

ip [ospf < adj | event | flood | lsa-generation | packet [packet-type] | retransmission | spf >]

For the configured debug destination(s):

ospf < adj | event | flood | lsa-generation | packet [packet-type] | retransmission | spf > — *Enables the specified IP-OSPF message type.*

adj — *Adjacency changes.*

event — *OSPF events.*

flood — *Information on flood messages.*

lsa-generation — *New LSAs added to database.*

packet [**packet-type**] — *All OSPF packet messages sent and received on the switch, where **packet-type** enables only the specified OSPF packet type. Valid values are:*

dd — *Database descriptions*

hello — *Hello messages*

lsa — *Link-state advertisements*

lsr — *Link-state requests*

lsu — *Link-state updates*

retransmission — *Retransmission timer messages.*

spf — *Path recalculation messages.*

ip [rip < database | event | trigger >]

rip < database | event | trigger > — *Enables the specified RIP message type for the configured destination(s).*

database — *Display database changes.*

event — *Display RIP events.*

trigger — *Display trigger messages.*

ipv6
[dhcpv6-client [events | packet]]
[forwarding | nd | packet]

When no debug options are included, displays debug messages for all IPv6 debug options.

dhcpv6-client [events | packet]: *Displays DHCPv6 client event and packet data.*

[forwarding]: *Displays IPv6 forwarding messages.*

[nd]: *Displays debug messages for IPv6 neighbor discovery.*

[packet]: *Displays IPv6 packet messages.*

lldp

Enables all LLDP message types for the configured destinations.

Debug Destinations

Use the **debug destination** command to enable (and disable) Syslog messaging on a Syslog server or to a CLI session for specified types of debug and Event Log messages.

Syntax: [no] debug destination < logging | session | buffer | debug-console>

logging

*Enables Syslog logging to configured Syslog servers so that the debug message types specified by the **debug <debug-type>** command (see “Debug Messages” on page C-50) are sent. (Default: Logging disabled)*

To configure a Syslog server IP address, refer to “Configuring a Syslog Server” on page C-55.

Note: *Debug messages from the switches covered in this guide have a debug severity level. Because the default configuration of some Syslog servers ignore Syslog messages with the debug severity level, ensure that the Syslog servers you want to use to receive debug messages are configured to accept the debug level. For more information, refer to “Operating Notes for Debug and Syslog” on page C-60.*

session

*Enables transmission of event notification messages to the CLI session that most recently executed this command. The session can be on any one terminal emulation device with serial, Telnet, or SSH access to the CLI at the Manager level prompt (**ProCurve#_**). If more than one terminal device has a console session with the CLI, you can redirect the destination from the current device to another device. Do so by executing **debug destination session** in the CLI on the terminal device on which you now want to display event messages.*

*Event message types received on the selected CLI session are configured with the **debug < debug-type >** command. (Refer to “Debug Messages” on page C-50.)*

buffer

*Enables Syslog logging to send the debug message types specified by the **debug < debug-type >** command to a buffer in switch memory. To view the debug messages stored in the switch buffer, enter the **show debug buffer** command.*

Logging Command

At the global configuration level, the **logging** command allows you to enable debug logging on specified Syslog servers and select a subset of Event Log messages to send for debugging purposes according to:

- Severity level
- System module

By specifying both a severity level and system module, you can use both configured settings to filter the Event Log messages you want to use to troubleshoot switch or network error conditions.

Caution

After you configure a Syslog server and a severity level and/or system module to filter the Event Log messages that are sent, if you save these settings to the startup configuration file by entering the **write memory** command, these debug and logging settings are automatically re-activated after a switch reboot or power recycle. The debug settings and destinations configured in your previous troubleshooting session will then be applied to the current session, which may not be desirable.

After a reboot, messages remain in the Event Log and are not deleted. However, after a power recycle, all Event Log messages are deleted.

If you configure a severity level and/or system module to temporarily filter Event Log messages, be sure to reset the values to their default settings by entering the **no** form of the following commands to ensure that Event Log messages of all severity levels and from all system modules are sent to configured Syslog servers:

```
ProCurve(config)# no logging severity <debug|major|error|warning|info>  
ProCurve(config)# no logging system-module <system-module >
```

Configuring a Syslog Server

Syslog is a client-server logging tool that allows a client switch to send event notification messages to a networked device operating with Syslog server software. Messages sent to a Syslog server can be stored to a file for later debugging analysis.

To use the Syslog feature, you must install and configure a Syslog server application on a networked host accessible to the switch. Refer to the documentation for the Syslog server application for instructions.

To configure a Syslog server, use the **logging** *< syslog-ip-addr >* command as described below.

When you configure a Syslog server, Event Log messages are automatically enabled to be sent to the server. To reconfigure this setting, use the following commands:

- Use **debug** command to specify additional debug message types (see “Debug Messages” on page C-50).
- Use the **logging** command to configure the system module or severity level used to filter the Event Log messages sent to configured Syslog servers (see “Configuring the Severity Level for Event Log Messages Sent to a Syslog Server” on page C-59 and “Configuring the System Module Used to Select the Event Log Messages Sent to a Syslog Server” on page C-60).

To display the currently configured Syslog servers as well as the types of debug messages and the severity-level and system-module filters used to specify the Event Log messages that are sent, enter the **show debug** command (see “Displaying a Debug/Syslog Configuration” on page C-46).

Syntax: [no] logging < syslog-ip-addr >

Enables or disables Syslog messaging to the specified IP address. You can configure up to six addresses. If you configure an address when none are already configured, this command enables destination logging (Syslog) and the Event debug type. Therefore, at a minimum, the switch begins sending Event Log messages to configured Syslog servers. The ACL, IP-OSPF, and/or IP-RIP message types will also be sent to the Syslog server(s) if they are currently enabled as debug types. (Refer to “Debug Messages” on page C-50.)

no logging removes all currently configured Syslog logging destinations from the running configuration.

no logging < syslog-ip-address > removes only the specified Syslog logging destination from the running configuration.

If you use the “no” form of the command to delete the only remaining Syslog server address, debug destination logging is disabled on the switch, but the default Event debug type is not changed.

*Also, removing all configured Syslog destinations with the **no logging** command (or a specified Syslog server destination with the **no logging < syslog-ip-address >** command) does not delete the Syslog server IP addresses stored in the startup configuration. To delete Syslog addresses in the startup configuration, you must enter a **no logging** command followed by the **write memory** command. To verify the deletion of a Syslog server address, display the startup configuration by entering the **show config** command.*

*To block the messages sent to configured Syslog servers from the currently configured debug message type, enter the **no debug < debug-type >** command. (See “Debug Messages” on page C-50.)*

*To disable Syslog logging on the switch without deleting configured server addresses, enter the **no debug destination logging** command. Note that, unlike the case in which no Syslog servers are configured, if one or more Syslog servers are already configured and Syslog messaging is disabled, configuring a new server address does not re-enable Syslog messaging. To re-enable Syslog messaging, you must enter the **debug destination logging** command.*

Syntax: [no] logging facility < facility-name >

The logging facility specifies the destination subsystem used in a configured Syslog server. (All configured Syslog servers must use the same subsystem.) ProCurve recommends the default (user) subsystem unless your application specifically requires another subsystem. Options include:

user (default) — Random user-level messages
kern — Kernel messages
mail — Mail system
daemon — System daemons
auth — Security/Authorization messages
syslog — Messages generated internally by Syslog
lpr — Line-Printer subsystem
news — Netnews subsystem
uucp — uucp subsystem
cron — cron/at subsystem
sys9 — cron/at subsystem
sys10 - sys14 — Reserved for system use
local10 - local17 — Reserved for system use

*Use the **no** form of the command to remove the configured facility and reconfigure the default (**user**) value. For a list of supported ProCurve switches, refer to the Note on page C-43.*

Adding a Description for a Syslog Server

You can associate a user-friendly description with each of the IP addresses (IPv4 only) configured for syslog using the CLI or SNMP.

Note

The HP enterprise MIB `hpicfSyslog.mib` allows the configuration and monitoring of syslog for SNMP (RFC 3164 supported).

The CLI command is:

Syntax: logging <ip-addr> control-descr <text_string>
no logging <ip-addr> [control-descr]

*An optional user-friendly description that can be associated with a server IP address. If no description is entered, this is blank. If <text_string> contains white space, use quotes around the string. IPv4 addresses only. Use the **no** form of the command to remove the description.*

Limit: 255 characters

Note: To remove the description using SNMP, set the description to an empty string.

```
ProCurve(config)# logging 10.10.10.2 control-descr syslog_one
```

Figure C-19. Example of the Logging Command with a Control Description

Caution

Entering the **no logging** command removes ALL the syslog server addresses without a verification prompt.

Adding a Priority Description

You can add a user-friendly description for the set of syslog filter parameters using the **priority-descr** option. The description can be added with the CLI or SNMP. The CLI command is:

Syntax: logging priority-descr <text_string>
no logging priority-descr

*Provides a user-friendly description for the combined filter values of **severity** and **system module**. If no description is entered, this is blank. If <text_string> contains white space, use quotes around the string. Use the **no** form of the command to remove the description.*

Limit: 255 characters

```
ProCurve(config)# logging priority-descr severe-pri
```

Figure C-20. Example of the Logging Command with a Priority Description

Note

A notification is sent to the SNMP agent if there are any changes to the syslog parameters either through the CLI or with SNMP.

Configuring the Severity Level for Event Log Messages Sent to a Syslog Server

Event Log messages are entered with one of the following severity levels (from highest to lowest):

Major: A fatal error condition has occurred on the switch.

Error: An error condition has occurred on the switch.

Warning: A switch service has behaved unexpectedly.

Information: Information on a normal switch event.

Debug: Reserved for ProCurve internal diagnostic information.

Using the **logging severity** command, you can select a set of Event Log messages according to their severity level and send them to a Syslog server. Messages of the selected and higher severity will be sent. To configure a Syslog server, see “Configuring a Syslog Server” on page C-55.

Syntax: [no] logging severity < major | error | warning | info | debug >

Configures the switch to send all Event Log messages with a severity level equal to or higher than the specified value to all configured Syslog servers.

*Default: **debug** (Reports messages of all severity levels.)*

*Use the **no** form of the command to remove the configured severity level and reconfigure the default value, which sends Event Log messages of all severity levels to Syslog servers.*

Note: *The severity setting does not affect event notification messages that the switch normally sends to the Event Log. All messages remain recorded in the Event Log.*

Configuring the System Module Used to Select the Event Log Messages Sent to a Syslog Server

Event Log messages contain the name of the system module that reported the event. Using the **logging system-module** command, you can select a set of Event Log messages according to the originating system module and send them to a Syslog server. To configure a Syslog server, see “Configuring a Syslog Server” on page C-55.

Using the **logging system-module** command, you can select messages from only one system module to be sent to a Syslog server. You cannot configure messages from multiple system modules to be sent. If you re-enter the command with a different system module name, the currently configured value is replaced with the new one.

Syntax: [no] logging system-module < system-module >

Configures the switch to send all Event Log messages being logged from the specified system module to configured Syslog servers.

Refer to Table C-1 on page C-27 for the correct value to enter for each system module.

*Default: **all-pass** (Reports all Event Log messages.)*

*Use the **no** form of the command to remove the configured system module value and reconfigure the default value, which sends Event Log messages from all system modules to Syslog servers.*

Note: *This setting has no effect on event notification messages that the switch normally sends to the Event Log.*

Operating Notes for Debug and Syslog

- **Rebooting the Switch or pressing the Reset button resets the Debug Configuration.**

Debug Option	Effect of a Reboot or Reset
logging (debug destination)	If Syslog server IP addresses are stored in the startup-config file, they are saved across a reboot and the logging destination option remains enabled. Otherwise, the logging destination is disabled.
session (debug destination)	Disabled.

Debug Option	Effect of a Reboot or Reset
ACL (debug type)	Disabled.
All (debug type)	Disabled.
event (debug type)	If a Syslog server IP address is configured in the startup-config file, the sending of Event Log messages is reset to enabled , regardless of the last active setting. If no Syslog server is configured, the sending of Event Log messages is disabled .
IP (debug type)	Disabled.

- **Debug commands do not affect normal message output to the Event Log.**

Using the **debug event** command, you can specify that Event Log messages are sent to the debug destinations you configure (CLI session and/or Syslog servers) in addition to the Event Log.

- **Ensure that your Syslog servers accept Debug messages.**

All Syslog messages resulting from a debug operation have a “debug” severity level. If you configure the switch to send debug messages to a Syslog server, ensure that the server’s Syslog application is configured to accept the “debug” severity level. (The default configuration for some Syslog applications ignores the “debug” severity level.)

- Duplicate IP addresses are not stored in the list of syslog servers.
- If the default severity value is in effect, all messages that have severities greater than the default value are passed to syslog. For example, if the default severity is “debug”, all messages that have severities greater than debug are passed to syslog.
- There is a limit of six syslog servers. All syslog servers are sent the same messages using the same filter parameters. An error is generated for an attempt to add more than six syslog servers.

Diagnostic Tools

Diagnostic Features

Feature	Default	Menu	CLI	Web
Port Auto negotiation	n/a	—	—	—
Ping test	n/a	—	page C-65	page C-64
Link test	n/a	—	page C-65	page C-64
Traceroute operation	n/a	—	page C-67	n/a
View switch configuration files	n/a	—	page C-71	page C-71
View switch (show tech) operation	n/a	—	page C-72	—
View crash information and command history	n/a	—	page C-78	—
View system information and software version	n/a	—	page C-78	—
Useful commands in a troubleshooting session	n/a	—	page C-82	—
Resetting factory-default configuration	page C-83 (Buttons)	—	page C-83	—
Restoring a flash image	n/a	—	page C-84	—
Port Status	n/a	page B-14	page B-14	page B-14

Port Auto-Negotiation

When a link LED does not light (indicating loss of link between two devices), the most common reason is a failure of port auto-negotiation between the connecting ports. If a link LED fails to light when you connect the switch to a port on another device, do the following:

1. Ensure that the switch port and the port on the attached end-node are both set to **Auto** mode.
2. If the attached end-node does not have an **Auto** mode setting, then you must manually configure the switch port to the same setting as the end-node port. Refer to Chapter 10, “Port Status and Configuration”.

Ping and Link Tests

The Ping test and the Link test are point-to-point tests between your switch and another IEEE 802.3-compliant device on your network. These tests can tell you whether the switch is communicating properly with another device.

Note

To respond to a Ping test or a Link test, the device you are trying to reach must be IEEE 802.3-compliant.

Ping Test. This is a test of the path between the switch and another device on the same or another IP network that can respond to IP packets (ICMP Echo Requests). To use the **ping** (or **tracert**) command with host names or fully qualified domain names, refer to “DNS Resolver” on page C-87.

Link Test. This is a test of the connection between the switch and a designated network device on the same LAN (or VLAN, if configured). During the link test, IEEE 802.2 test packets are sent to the designated network device in the same VLAN or broadcast domain. The remote device must be able to respond with an 802.2 Test Response Packet.

Web: Executing Ping or Link Tests

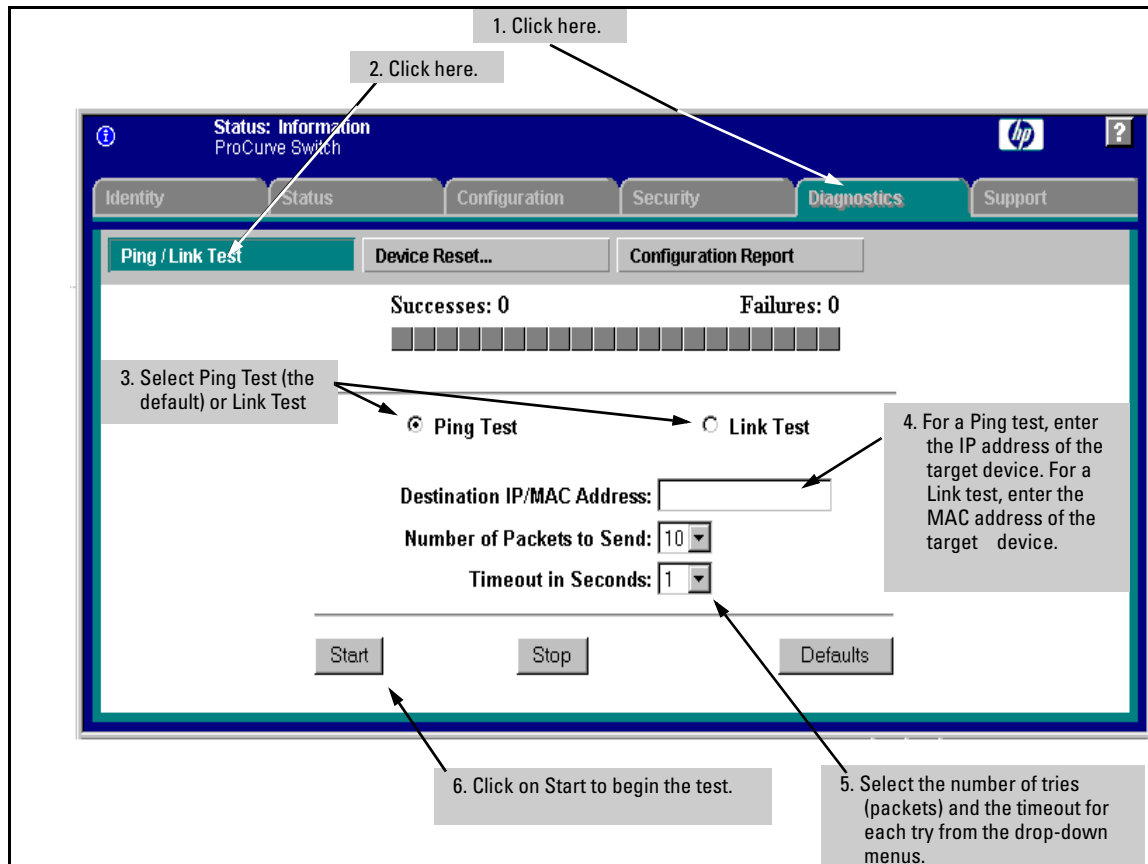


Figure C-21. Link and Ping Test Screen on the Web Browser Interface

Successes indicates the number of Ping or Link packets that successfully completed the most recent test.

Failures indicates the number of Ping or Link packets that were unsuccessful in the last test. Failures indicate connectivity or network performance problems (such as overloaded links or devices).

Destination IP/MAC Address is the network address of the target, or destination, device to which you want to test a connection with the switch. An IP address is in the X.X.X.X format where X is a decimal number between 0 and 255. A MAC address is made up of 12 hexadecimal digits, for example, 0060b0-080400.

Number of Packets to Send is the number of times you want the switch to attempt to test a connection.

Timeout in Seconds is the number of seconds to allow per attempt to test a connection before determining that the current attempt has failed.

To halt a Link or Ping test before it concludes, click on the Stop button.

To reset the screen to its default settings, click on the Defaults button.

CLI: Ping Test

The Ping (Packet InterNet Groper) test uses Internet Control Message Protocol (ICMP) echo requests and ICMP echo replies to determine if another device is alive. It also measures the amount of time it takes to receive a reply from the specified destination. The Ping command has several extended commands that allow advanced checking of destination availability.

Syntax: ping <ip-address | hostname | switch-num> [repetitions <1-10000>]
[timeout <1-60>] [source <ip-address> | <vlan-id>] [data-size <0 - 65471>]
[data-fill <0-1024>]

ping6 <ip-address | hostname | switch-num> [repetitions <1-10000>]
[timeout <1-60>] [source <ip-address> | <vlan-id>] [data-size <0 - 65471>]
[data-fill <0-1024>]

Sends ICMP echo requests to determine if another device is alive.

Note: For information about **ping6**, see the “IPv6 Configuration Guide” for your switch.

<ip-address | hostname>

Target IP address or hostname of the destination node being pinged.

repetitions <1-10000>

*Number of ping packets sent to the destination address.
Default: 1*

timeout <1-60>

*Timeout interval in seconds; the ECHO REPLY must be received before this time interval expires for the Ping to be successful.
Default: 5*

source <ip-addr | hostname >

Source IP address or hostname. The source IP address must be owned by the router. If a VLAN is specified, the IP address associated with the specified VLAN is used.

data-size <0-65471>

Size of packet sent. Default: 0 (zero)

data-fill <0-1024>

The data pattern in the packet. Default: Zero length string

Basic Ping Operation	→	ProCurve > ping 10.28.227.103 10.28.227.103 is alive, time = 15 ms
Ping with Repetitions	→	ProCurve > ping 10.28.227.103 repetitions 3 10.28.227.103 is alive, iteration 1, time = 15 ms 10.28.227.103 is alive, iteration 2, time = 15 ms 10.28.227.103 is alive, iteration 3, time = 15 ms
Ping with Repetitions and Timeout	→	ProCurve > ping 10.28.227.103 repetitions 3 timeout 2 10.28.227.103 is alive, iteration 1, time = 15 ms 10.28.227.103 is alive, iteration 2, time = 10 ms 10.28.227.103 is alive, iteration 3, time = 15 ms
Ping Failure	↘	ProCurve > ping 10.28.227.105 Target did not respond.

Figure C-22. Examples of Ping Tests

To halt a ping test before it concludes, press **[Ctrl] [C]**.

Note

To use the **ping** (or **tracert**) command with host names or fully qualified domain names, refer to “DNS Resolver” on page C-87.

Link Tests

You can issue single or multiple link tests with varying repetitions and timeout periods. The defaults are:

- Repetitions: 1 (1 - 999)
- Timeout: 5 seconds (1 - 256 seconds)

Syntax: link < mac-address > [repetitions < 1 - 999 >] [timeout < 1 - 256 >]
[vlan < vlan-id >]

Basic Link Test	ProCurve# link 0030c1-7fcc40 Link-test passed.
Link Test with Repetitions	ProCurve# link 0030c1-7fcc40 repetitions 3 802.2 TEST packets sent: 3, responses received: 3
Link Test with Repetitions and Timeout	ProCurve# link 0030c1-7fcc40 repetitions 3 timeout 1 802.2 TEST packets sent: 3, responses received: 3
Link Test Over a Specific VLAN	ProCurve# link 0030c1-7fcc40 repetitions 3 timeout 1 vlan 1 802.2 TEST packets sent: 3, responses received: 3
Link Test Over a Specific VLAN; Test Fail	ProCurve# link 0030c1-7fcc40 repetitions 3 timeout 1 vlan 222 802.2 TEST packets sent: 3, responses received: 0

Figure C-23. Example of Link Tests

Traceroute Command

The **traceroute** command enables you to trace the route from the switch to a host address.

This command outputs information for each (router) hop between the switch and the destination address. Note that every time you execute **traceroute**, it uses the same default settings unless you specify otherwise for that instance of the command.

Syntax: traceroute < ip-address | hostname >
traceroute6 < ip-address | hostname >

*Lists the IP address or hostname of each hop in the route, plus the time in microseconds for the **traceroute** packet reply to the switch for each hop.*

*To halt an ongoing traceroute search, press the **[Ctrl] [C]** keys.*

Note: For information about **traceroute6**, see the “IPv6 Configuration Guide” for your switch.

<ip-address | hostname>

The IP address or hostname of the device to which to send the traceroute.

[minttl < 1-255 >]

*For the current instance of **traceroute**, changes the minimum number of hops allowed for each probe packet sent along the route. If **minttl** is greater than the actual number of hops, then the output includes only the hops at and above the **minttl** threshold. (The hops below the threshold are not listed.) If **minttl** matches the actual number of hops, only that hop is shown in the output. If **minttl** is less than the actual number of hops, then all hops are listed. For any instance of **traceroute**, if you want a **minttl** value other than the default, you must specify that value. (Default: 1)*

[maxttl < 1-255 >]

*For the current instance of **traceroute**, changes the maximum number of hops allowed for each probe packet sent along the route. If the destination address is further from the switch than **maxttl** allows, then **traceroute** lists the IP addresses for all hops it detects up to the **maxttl** limit. For any instance of **traceroute**, if you want a **maxttl** value other than the default, you must specify that value. (Default: 30)*

[timeout < 1-120 >]

*For the current instance of **traceroute**, changes the timeout period the switch waits for each probe of a hop in the route. For any instance of **traceroute**, if you want a **timeout** value other than the default, you must specify that value. (Default: 5 seconds)*

[probes < 1-5 >]

*For the current instance of **traceroute**, changes the number of queries the switch sends for each hop in the route. For any instance of **traceroute**, if you want a **probes** value other than the default, you must specify that value. (Default: 3)*

[source <ip-addr | vlan-id>]

The source IP address or VLAN. The source IP address must be owned by the router. If a VLAN is specified, the IP address associated with the specified VLAN is used

A Low Maxttl Causes Traceroute To Halt Before Reaching the Destination Address. For example, executing **traceroute** with its default values for a destination IP address that is four hops away produces a result similar to this:


```

ProCurve# traceroute 125.25.24.35
traceroute to 125.25.24.35 ,
          1 hop min, 30 hops max, 5 sec. timeout, 3 probes
 1 10.255.120.2          0 ms      0 ms      0 ms
 2 10.71.217.2           7 ms      3 ms      0 ms
 3 10.243.170.1         0 ms      1 ms      0 ms
 4 125.25.24.35         3 ms      3 ms      0 ms

```

Intermediate router hops with the time taken for the switch to receive an acknowledgement of each probe reaching each router.

Destination IP Address

Figure C-24. Example of a Completed Traceroute Enquiry

Continuing from the previous example (Figure C-24, above), executing **traceroute** with an insufficient **maxttl** for the actual hop count produces an output similar to this:

```

Traceroute does not reach destination IP address because of low maxttl setting.
ProCurve# traceroute 125.25.24.35 (maxttl 3)
traceroute to 125.25.24.35 ,
          1 hop min, 3 hops max, 5 sec. timeout, 3 probes
 1 10.255.120.2          0 ms      0 ms      0 ms
 2 10.71.217.2           0 ms      0 ms      0 ms
 3 10.243.170.1         0 ms      *          0 ms

```

The asterisk indicates there was a timeout on the second probe to the third hop.

Figure C-25. Example of Incomplete Traceroute Due to Low Maxttl Setting

If A Network Condition Prevents Traceroute from Reaching the Destination. Common reasons for Traceroute failing to reach a destination include:

- Timeouts (indicated by one asterisk per probe, per hop; refer to Figure C-25, above.)
- Unreachable hosts
- Unreachable networks
- Interference from firewalls
- Hosts configured to avoid responding

Executing traceroute where the route becomes blocked or otherwise fails results in an output marked by timeouts for all probes beyond the last detected hop. For example with a maximum hop count of 7 (**maxttl = 7**), where the route becomes blocked or otherwise fails, the output appears similar to this:

```
ProCurve# traceroute 107.64.197.100 maxttl 7
traceroute to 107.64.197.100 ,
          1 hop min, 7 hops max, 5 sec. timeout, 3 probes
 1 10.255.120.2          0 ms      0 ms      0 ms
 2 10.71.217.2          0 ms      0 ms      0 ms
 3 * (10.243.170.1) *
 4 * * *
 5 * * *
 6 * * *
 7 * * *
```

At hop 3, the first and third probes timed out but the second probe reached the router. All further probes within the **maxttl** timed-out without finding a router or the destination IP address.

An asterisk indicates a timeout without finding the next hop.

Figure C-26. Example of Traceroute Failing to Reach the Destination Address

Viewing Switch Configuration and Operation

In some troubleshooting scenarios, you may need to view the switch configuration to diagnose a problem. The complete switch configuration is contained in a file that you can browse from either the web browser interface or the CLI using the commands described in this section.

CLI: Viewing the Startup or Running Configuration File

Using the CLI, you can display either the running or the startup configuration. For more information and examples of how to use these commands, refer to Chapter 6, “Switch Memory and Configuration”.)

Syntax: write terminal

Displays the running configuration.

show config

Displays the startup configuration.

show running-config

Displays the running-config file.

Web: Viewing the Configuration File

To display the running configuration, through the web browser interface:

1. Click on the **Diagnostics** tab.
2. Click on **[Configuration Report]**
3. Use the right-side scroll bar to scroll through the configuration listing.

CLI: Viewing a Summary of Switch Operational Data

Syntax: show tech

By default, the **show tech** command displays a single output of switch operating and running-configuration data from several internal switch sources, including:

- Image stamp (software version data)
- Running configuration
- Event Log listing
- Boot History
- Port settings
- Status and counters — port status
- IP routes
- Status and counters — VLAN information
- GVRP support
- Load balancing (trunk and LACP)

Figure C-27 shows sample output from the **show tech** command.

```
ProCurve# show tech

show system

Status and Counters - General System Information

System Name       : 5400_1
System Contact    :
System Location   :

MAC Age Time (sec) : 300

Time Zone         : 0
Daylight Time Rule : None

Software revision : K.14.XX           Base MAC Addr      : 001871-c42f00
ROM Version       : K.12.12           Serial Number     : SG641SU00L

Up Time          : 23 hours           Memory - Total    :
CPU Util (%)     : 10                  Free              :

IP Mgmt  - Pkts Rx : 759              Packet - Total    : 6750
          Pkts Tx  : 2                  Buffers  Free    : 5086
                                          Lowest   : 4961
                                          Missed   : 0

show flash
Image          Size(Bytes)  Date   Version  Build #
-----
-----
```

Figure C-27. Example of Show Tech Command

To specify the data displayed by the **show tech** command, use the **copy show tech** command as described in “Customizing show tech Command Output” on page C-75.

Saving show tech Command Output to a Text File

When you enter the **show tech** command, a summary of switch operational data is sent to your terminal emulator. You can use your terminal emulator’s text capture features to save the **show tech** data to a text file for viewing, printing, or sending to an associate to diagnose a problem.

For example, if your terminal emulator is the Hyperterminal application available with Microsoft® Windows® software, you can copy the **show tech** output to a file and then use either Microsoft Word or Notepad to display the data. (In this case, Microsoft Word provides the data in an easier-to-read format.)

The following example uses the Microsoft Windows terminal emulator. If you are using a different terminal emulator application, refer to the documentation provided with the application.

To save **show tech** command output from your terminal emulator to a text file, follow these steps:

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




Figure C-28. Capture Text window of the Hyperterminal Application

2. In the **File** field, enter the path and file name in which you want to store the **show tech** output.




Figure C-29. Entering a Path and Filename for Saving show tech Output

3. Click [**Start**] to create and open the text file.
4. From the global configuration context, enter the **show tech** command:

```
ProCurve# show tech
```

The **show tech** command output is copied into the text file and displayed on the terminal emulator screen. When the command output stops and displays -- MORE --, press the Space bar to display and copy more information. The CLI prompt appears when the command output finishes.

5. Click on **Transfer | Capture Text | Stop** in HyperTerminal to stop copying data and save the text file.

If you do not stop HyperTerminal from copying command output into the text file, additional unwanted data can be copied from the HyperTerminal screen.

6. To access the file, open it in Microsoft Word, Notepad, or a similar text editor.

Customizing show tech Command Output

Use the **copy show tech** command to customize the detailed switch information displayed with the **show tech** command to suit your troubleshooting needs.

To customize the information displayed with the **show tech** command:

1. Determine the information that you want to gather to troubleshoot a problem in switch operation.
2. Enter the **copy show tech** command to specify the data files that contain the information you want to view.

Syntax: copy <source> show- tech

*Specifies the operational and configuration data from one or more source files to be displayed by the **show tech** command. Enter the command once for each data file that you want to include in the display.*

Default: Displays data from all source files, where <source> can be any one of the following values:

command-output "< command >"

*Includes the output of a specified command in **show-tech** command output. Enter the command name between double-quotation marks; for example, **copy "show system" show-tech**.*

crash-data [slot-id | master]:

*Includes the crash data from all management and interface modules in **show tech** command output.*

To limit the amount of crash data displayed, specify an installed module or management modules, where:

slot-id: *Includes the crash data from an installed module. Valid slot IDs are the letters **a** through **h**.*

master: *Includes the crash data from both management modules.*

Syntax: copy <source> show- tech

crash-log [*slot-id* | master]:

*Includes the crash logs from all management and interface modules in **show tech** command output.*

To limit the amount of crash-log data displayed, specify an installed module or management modules, where:

slot-id: *Includes the crash log from an installed module.*

*Valid slot IDs are the letters **a** through **h**.*

master: *Includes the crash log from both management modules.*

event-log

*Copies the contents of the Event Log to **show tech** command output.*

running-config

*Includes the contents of the running configuration file in **show tech** command output.*

startup-config

*Includes the contents of the startup configuration file in **show tech** command output.*

tftp config < startup-config | running-config > < ip-addr > < remote-file >
< pc | unix >

*Downloads the contents of a configuration file from a remote host to **show tech** command output, where:*

ip-addr: *Specifies the IP address of the remote host device.*

remote-file: *Specifies the pathname on the remote host for the configuration file whose contents you want to include in the command output.*

pc | unix: *Specifies whether the remote host is a DOS-based PC or UNIX workstation.*

*For more information on using **copy tftp** commands, refer to the “File Transfers” appendix.*

Syntax: copy <source> show- tech

usb config < startup-config < filename > | command-file < acl-
filename.txt >

*Copies the contents of a configuration file or ACL command file from a USB flash drive to **show tech** command output, where:*

startup-config <filename >: Specifies the name of a startup configuration file on the USB drive.

command-file <acl-filename.txt >: Specifies the name of an ACL command file on the USB drive.

*For more information on using **copy usb** commands, refer to the “File Transfers” appendix.*

xmodem config < startup-config | config < filename > | command-file
< acl-filename.txt > < pc | unix >

*Copies the contents of a configuration file or ACL command file from a serially connected PC or UNIX workstation to **show tech** command output, where:*

startup-config: Specifies the name of the startup configuration file on the connected device.

config <filename >: Specifies the pathname of a configuration file on the connected device.

command-file <acl-filename.txt >: Specifies the pathname of an ACL command file on the connected device.

pc | unix: Specifies whether the connected device is a DOS-based PC or UNIX workstation.

*For more information on using **copy xmodem** commands, refer to the “File Transfers” appendix.*

CLI: Viewing More Information on Switch Operation

Use the following commands to display additional information on switch operation for troubleshooting purposes.

Syntax: show boot-history

Displays the crash information saved for each management module on the switch (see “Displaying Saved Crash Information” in the “Redundancy (Switch 8212z1)” chapter).

show history

Displays the current command history. This command output is used for reference or when you want to repeat a command (see “CLI: Useful Commands for Troubleshooting Sessions” on page C-82)

show system-information

Displays globally configured parameters and information on switch operation (see “CLI: Viewing and Configuring System Information” in the “Interface Access and System Information” chapter).

show version

Displays the software version currently running on the switch, and the flash image from which the switch booted (primary or secondary). For more information, see “Displaying Management Information” in the “Redundancy (Switch 8212z1)” chapter.

show interfaces

Displays information on the activity on all switch ports (see “CLI: Viewing Port Status and Configuring Port Parameters” in the “Port Status and Configuration” chapter).

show interfaces-display

*Displays the same information as the **show interfaces** command and dynamically updates the output every three seconds. Press Ctrl + C to stop the dynamic updates of system information. Use the Arrow keys to view information that is off the screen.*

Pattern Matching When Using the Show Command

The pattern matching option with the **show** command provides the ability to do searches for specific text. Selected portions of the output are displayed depending on the parameters chosen.

Syntax: show <command option> | <include | exclude | begin > <regular expression>

*Use matching pattern searches to display selected portions of the output from a **show** command. There is no limit to the number of characters that can be matched. Only regular expressions are permitted; symbols such as the asterisk cannot be substituted to perform more general matching.*

include Only the lines that contain the matching pattern are displayed in the output.

exclude: Only the lines that contain the matching pattern are not displayed in the output.

begin: The display of the output begins with the line that contains the matching pattern.

Note

Pattern matching is case-sensitive.

Below are examples of what portions of the running config file display depending on the option chosen.

```
ProCurve(config)# show run | include ipv6
  ipv6 enable
  ipv6 enable
ipv6 access-list "EH-01"
ProCurve(config)#
```

Displays only lines that contain "ipv6".

Figure C-30. Example of Pattern Matching with Include Option

```
ProCurve(config)# show run | exclude ipv6

Running configuration:

; J8697A Configuration Editor; Created on release #K.14.06

hostname "ProCurve Switch 5406z1"
module 1 type J8702A
module 2 type J8705A
snmp-server community "notpublic" Unrestricted
vlan 1
    name "DEFAULT_VLAN"
    untagged A1-A24,B1-B20
    ip address dhcp-bootp
    no untagged B21-B24
    exit
vlan 20
    name "VLAN20"
    untagged B21-B24
    no ip address
    exit
policy qos "michael"
    exit
    sequence 10 deny tcp 2001:db8:255::/48 2001:db8:125::/48
    exit
no autorun
password manager

ProCurve(config)#
```

Displays all lines that don't contain "ipv6".

Figure C-31. Example of Pattern Matching with Exclude Option

```
ProCurve(config)# show run | begin ipv6
  ipv6 enable
  no untagged B21-B24
  exit
vlan 20
  name "VLAN20"
  untagged B21-B24
  ipv6 enable
  no ip address
  exit
policy qos "michael"
  exit
ipv6 access-list "EH-01"
  sequence 10 deny tcp 2001:db8:255::/48 2001:db8:125::/48
  exit
no autorun
password manager

ProCurve(config)#
```

Displays the running config beginning at the first line that contains "ipv6".

Figure C-32. Example of Pattern Matching with Begin Option

Figure C-33 is an example of the **show arp** command output, and then the output displayed when the **include** option has the IP address of **15.255.128.1** as the regular expression.

```
ProCurve(config)# show arp

IP ARP table

  IP Address          MAC Address          Type      Port
  -----
  15.255.128.1       00000c-07ac00       dynamic  B1
  15.255.131.19     00a0c9-b1503d       dynamic
  15.255.133.150    000bcd-3cbeec       dynamic  B1

ProCurve(config)# show arp | include 15.255.128.1
  15.255.128.1       00000c-07ac00       dynamic  B1
```

Figure C-33. Example of the Show ARP Command and Pattern Matching with the Include Option

CLI: Useful Commands for Troubleshooting Sessions

Use the following commands in a troubleshooting session to more accurately display the information you need to diagnose a problem. For more information on other these CLI practices, refer to chapter 4, “Using the Command Line Interface (CLI)”.

Syntax: alias

Creates a shortcut alias name for commonly used commands and command options.

For more information, see “Using a Command Alias” in the “Using the Command Line Interface (CLI)” chapter.

kill

*Terminates a currently running, remote troubleshooting session. Use the **show ip ssh command** to list the current management sessions.*

For more information, see “Denying Interface Access by Terminating Remote Management Sessions” in the “Interface Access and System Information” chapter.

[no] page

*Toggles the paging mode for **show** commands between continuous listing and per-page listing.*

repeat

Repeatedly executes one or more commands so that you can see the results of multiple commands displayed over a period of time. To halt the command execution, press any key on the keyboard.

For more information, see “Repeating a Command” in the “Using the Command Line Interface (CLI)” chapter.

setup

Displays the Switch Setup screen from the menu interface.

Restoring the Factory-Default Configuration

As part of your troubleshooting process, it may become necessary to return the switch configuration to the factory default settings. This process momentarily interrupts the switch operation, clears any passwords, clears the console Event Log, resets the network counters to zero, performs a complete self test, and reboots the switch into its factory default configuration including deleting an IP address. There are two methods for resetting to the factory-default configuration:

- CLI
- Clear/Reset button combination

Note

ProCurve recommends that you save your configuration to a TFTP server before resetting the switch to its factory-default configuration. You can also save your configuration via Xmodem, to a directly connected PC.

CLI: Resetting to the Factory-Default Configuration

This command operates at any level *except* the Operator level.

Syntax: erase startup-configuration

Deletes the startup-config file in flash so that the switch will reboot with its factory-default configuration.

Note

The **erase startup-config** command does not clear passwords.

Clear/Reset: Resetting to the Factory-Default Configuration

To execute the factory default reset, perform these steps:

1. Using pointed objects, simultaneously press both the Reset and Clear buttons on the front of the switch.

2. Continue to press the Clear button while releasing the Reset button.
3. When the Self Test LED begins to flash, release the Clear button.

The switch will then complete its self test and begin operating with the configuration restored to the factory default settings.

Restoring a Flash Image

The switch can lose its operating system if either the primary or secondary flash image location is empty or contains a corrupted OS file and an operator uses the **erase flash** command to erase a good OS image file from the opposite flash location.

To Recover from an Empty or Corrupted Flash State. Use the switch's console serial port to connect to a workstation or laptop computer that has the following:

- A terminal emulator program with Xmodem capability, such as the HyperTerminal program included in Windows PC software.
- A copy of a good OS image file for the switch.

Note

The following procedure requires the use of Xmodem, and copies an OS image into primary flash only.

This procedure assumes you are using HyperTerminal as your terminal emulator. If you use a different terminal emulator, you may need to adapt this procedure to the operation of your particular emulator.

1. Start the terminal emulator program.
2. Ensure that the terminal program is configured as follows:
 - Baud rate: 9600
 - 1 stop bit
 - No parity
 - No flow control
 - 8 Bits
3. Use the Reset button to reset the switch. The following prompt should then appear in the terminal emulator:

Enter h or ? for help.

=>

4. Since the OS file is large, you can increase the speed of the download by changing the switch console and terminal emulator baud rates to a high speed. For example:

- a. Change the switch baud rate to 115,200 Bps.

=> sp 115200

- b. Change the terminal emulator baud rate to match the switch speed:
 - i. In HyperTerminal, select **Call | Disconnect**.
 - ii. Select **File | Properties**.
 - iii. Click on **Configure**.
 - iv. Change the baud rate to **115200**.
 - v. Click on **[OK]**. In the next window, click on **[OK]** again.
 - vi. Select **Call | Connect**
 - vii. Press **[Enter]** one or more times to display the => prompt.

5. Start the Console Download utility by typing **do** at the => prompt and pressing **[Enter]**:

=> do

6. You will then see this prompt:

```
You have invoked the console download utility.  
Do you wish to continue? (Y/N)>_
```

7. At the above prompt:
 - a. Type **Y** (for Yes)
 - b. Select **Transfer | File** in HyperTerminal.
 - c. Enter the appropriate filename and path for the OS image.
 - d. Select the **Xmodem** protocol (and not the 1k Xmodem protocol).
 - e. Click on **[Send]**.

If you are using HyperTerminal, you will see a screen similar to the following to indicate that the download is in progress:




Figure C-34. Example of Xmodem Download in Progress

8. When the download completes, the switch reboots from primary flash using the OS image you downloaded in the preceding steps, plus the most recent startup-config file.

DNS Resolver

The Domain Name System (DNS) resolver is designed for use in local network domains where it enables use of a host name or fully qualified domain name with DNS-compatible switch CLI commands. (At software release K.13.01, the DNS-compatible commands include **ping** and **traceroute**.)

Beginning with software release K.13.01, DNS operation supports both IPv4 and IPv6 DNS resolution and multiple, prioritized DNS servers. (For information on IPv6 DNS resolution, refer to the latest *IPv6 Configuration Guide* for your switch.)

Terminology

Domain Suffix — Includes all labels to the right of the unique host name in a fully qualified domain name assigned to an IP address. For example, in the fully qualified domain name “device53.evergreen.trees.org”, the domain suffix is “evergreen.trees.org”, while “device53” is the unique (host) name assigned to a specific IP address.

Fully Qualified Domain Name — The sequence of labels in a domain name identifying a specific host (host name) and the domain in which it exists. For example, if a device with an IP address of 10.10.10.101 has a host name of *device53* and resides in the *evergreen.trees.org* domain, then the device’s fully qualified domain name is *device53.evergreen.trees.org* and the DNS resolution of this name is 10.10.10.101.

Host Name — The unique, leftmost label in a domain name assigned to a specific IP address in a DNS server configuration. This enables the server to distinguish a device using that IP address from other devices in the same domain. For example, in the *evergreen.trees.org* domain, if an IPv4 address of 10.10.100.27 is assigned a host name of *accounts015* and another IP address of 10.10.100.33 is assigned a host name of *sales021*, then the switch configured with the domain suffix *evergreen.trees.org* and a DNS server that resolves addresses in that domain can use the host names to reach the devices with DNS-compatible commands. For example:

```
ping accounts015
traceroute accounts015
```

Basic Operation

- When the switch is configured with only the IP address of a DNS server available to the switch, then a DNS-compatible command, executed with a fully qualified domain name, can reach a device found in any domain accessible through the configured DNS server.
- When the switch is configured with both of the following:
 - the IP address of a DNS server available to the switch
 - the domain suffix of a domain available to the configured DNS server

then:

- A DNS-compatible command that includes the host name of a device in the same domain as the configured domain suffix can reach that device.
- A DNS-compatible command that includes a fully qualified domain name can reach a device in any domain that is available to the configured DNS server.

Example. Suppose the switch is configured with the domain suffix **mygroup.procurve.net** and the IP address for an accessible DNS server. If an operator wants to use the switch to ping a target host in this domain by using the DNS name “leader” (assigned by a DNS server to an IP address used in that domain), then the operator can use either of the following commands:

```
ProCurve# ping leader
10.28.229.220 is alive, time = 1 ms

ProCurve# ping leader.mygroup.procurve.net
10.28.229.220 is alive, time = 1 ms
```

Figure C-35. Example of Using Either a Host Name or a Fully Qualified Domain Name

In the preceding example, if the DNS server’s IP address is configured on the switch, but a domain suffix is either not configured or is configured for a different domain than the target host, then the fully qualified domain name *must* be used.

Note that if the target host is in a domain *other than* the domain configured on the switch, then:

- The host's domain must be reachable from the switch. This requires that the DNS server for the switch must be able to communicate with the DNS server(s) in the path to the domain in which the target host operates.
- The fully qualified domain name must be used, and the domain suffix must correspond to the domain in which the target host operates, regardless of the domain suffix configured in the switch.

Example. Suppose the switch is configured with the domain suffix **mygroup.procurve.net** and the IP address for an accessible DNS server in this same domain. This time, the operator wants to use the switch to trace the route to a host named "remote-01" in a different domain named **common.group.net**. Assuming this second domain is accessible to the DNS server already configured on the switch, a **traceroute** command using the target's fully qualified DNS name should succeed.

```
ProCurve# traceroute [remote-01.common.group.net]
[traceroute to 10.22.240.73]
      1 hop min, 30 hops max, 5 sec. timeout, 3 probes
 1 10.28.229.3          0 ms      0 ms      0 ms
 2 10.71.217.1         0 ms      0 ms      0 ms
 3 10.0.198.2          1 ms      0 ms      0 ms
[4 10.22.240.73       0 ms      0 ms      0 ms]
```

Fully Qualified Host Name for the Target Host

IP Address for Target Host "remote-01"

Figure C-36. Example Using the Fully Qualified Domain Name for an Accessible Target in Another Domain

Configuring and Using DNS Resolution with DNS-Compatible Commands

(At software release K.13.01, the DNS-compatible commands include **ping** and **traceroute**.)

1. Determine the following:
 - a. The IP address for a DNS server operating in a domain in your network
 - b. The priority (1 - 3) of the selected server, relative to other DNS servers in the domain

- c. The domain name for an accessible domain in which there are hosts you want to reach with a DNS-compatible command. (This is the domain suffix in the fully qualified domain name for a given host operating in the selected domain. Refer to “Terminology” on page C-87.) Note that if a domain suffix is not configured, fully qualified domain names can be used to resolve DNS-compatible commands.
 - d. the host names assigned to target IP addresses in the DNS server for the specified domain
2. Use the data from steps 1a through 1c to configure the DNS entry on the switch.
 3. Use a DNS-compatible command with the host name to reach the target devices.

Configuring a DNS Entry

The switch allows up to three DNS server entries (IP addresses for DNS servers). One domain suffix can also be configured to support resolution of DNS names in that domain by using a host name only. Including the domain suffix enables the use of DNS-compatible commands with a target’s host name instead of the target’s fully qualified domain name.

Syntax: [no] ip dns server-address priority < 1 - 3 > < ip-addr >

Configures the access priority and IP address of a DNS server accessible to the switch. These settings specify:

- *the relative priority of the DNS server when multiple servers are configured*
- *the IP address of the DNS server*

These settings must be configured before a DNS-compatible command can be executed with host name criteria.

*The switch supports three prioritized DNS server entries. Configuring another IP address for a priority that has already been assigned to an IP address is not allowed. To replace one IP address at a given priority level with another address having the same priority, you must first use the **no** form of the command to remove the unwanted address. Also, only one instance of a given server address is allowed in the server list. Attempting to enter a duplicate of an existing entry at a different priority level is not allowed. To change the priority of an existing server address, use the **no** form of the command to remove the entry, then re-enter the address with the new priority.*

*The **no** form of the command replaces the configured IP address with the null setting. (Default: null)*

Syntax: [no] ip dns domain-name < domain-name-suffix >

This optional DNS command configures the domain suffix that is automatically appended to the host name entered with a DNS-compatible command. When the domain suffix and the IP address for a DNS server that can access that domain are both configured on the switch, you can execute a DNS-compatible command using only the host name of the desired target. (For an example, refer to Figure C-35 on page C-88.) In either of the following two instances, you must manually provide the domain identification by using a fully qualified DNS name with a DNS-compatible command:

- *If the DNS server IP address is configured on the switch, but the domain suffix is not configured (null)*
- *The domain suffix configured on the switch is not the domain in which the target host exists*

The switch supports one domain suffix entry and three DNS server IP address entries. (Refer to the preceding command description.)

*The **no** form of the command replaces the configured domain suffix with the null setting. (Default: null)*

Example Using DNS Names with Ping and Traceroute

In the network illustrated in Figure C-37, the switch at 10.28.192.1 is configured to use DNS names for DNS-compatible commands in the *pubs.outdoors.com* domain. The DNS server has been configured to assign the host name *docservr* to the IP address used by the document server (10.28.229.219).

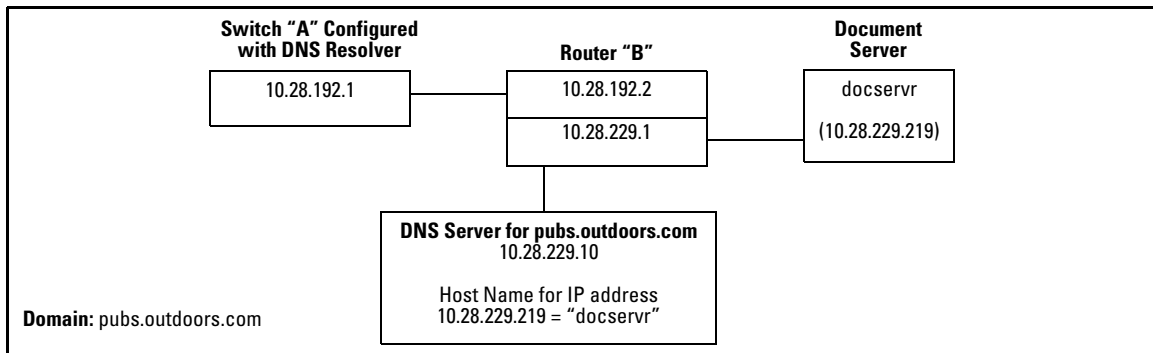


Figure C-37. Example Network Domain

Configuring switch “A” with the domain name and the IP address of a DNS server for the domain enables the switch to use host names assigned to IP addresses in the domain to perform **ping** and **traceroute** actions on the devices in the domain. To summarize:

Entity:	Identity:
DNS Server IP Address	10.28.229.10
Domain Name (and Domain Suffix for Hosts in the Domain)	pubs.outdoors.com
Host Name Assigned to 10.28.229.219 by the DNS Server	docservr
Fully Qualified Domain Name for the IP address Used By the Document Server (10.28.229.219)	docservr.pubs.outdoors.com
Switch IP Address	10.28.192.1
Document Server IP Address	10.28.229.219

With the above already configured, the following commands enable a DNS-compatible command with the host name **docservr** to reach the document server at 10.28.229.219.

```
ProCurve(config)# ip dns server-address 10.28.229.10
ProCurve(config)# ip dns domain-name pubs.outdoors.com
```

Figure C-38. Configuring Switch “A” in FigureC-37 To Support DNS Resolution

```
ProCurve# ping docservr
10.28.229.219 is alive, time = 1 ms

ProCurve# traceroute docservr
traceroute to 10.28.229.219
    1 hop min, 30 hops max, 5 sec. timeout, 3 probes
 1 10.28.192.2    1 ms    0 ms    0 ms
 2 10.28.229.219  0 ms    0 ms    0 ms
```

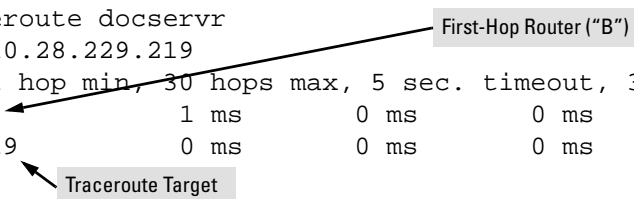


Figure C-39. Example of Ping and Traceroute Execution for the Network in Figure C-37 on Page C-91

As mentioned under “Basic Operation” on page C-88, if the DNS entry configured in the switch does not include the domain suffix for the desired target, then you must use the target host’s fully qualified domain name with DNS-compatible commands. For example, using the document server in Figure C-37 as a target:

```
ProCurve# ping [docsrvr.pubs.outdoors.com]
10.28.229.219 is alive, time = 1 ms

ProCurve# traceroute docsrvr[.pubs.outdoors.com]
traceroute to 10.28.229.219
          1 hop min, 30 hops max, 5 sec. timeout, 3 probes
 1 10.28.192.2          1 ms          0 ms          0 ms
 2 10.28.229.219       0 ms          0 ms          0 ms
```

Target's Fully Qualified Domain Name

Figure C-40. Example of Ping and Traceroute Execution When Only the DNS Server IP Address Is Configured

Viewing the Current DNS Configuration

The **show ip** command displays the current domain suffix and the IP address of the highest priority DNS server configured on the switch, along with other IP configuration information. If the switch configuration currently includes a non-default (non-null) DNS entry, it will also appear in the **show run** command output.

```
ProCurve# show ip

Internet (IP) Service

IP Routing : Disabled

Default Gateway : 10.28.192.2
Default TTL     : 64
Arp Age        : 20
[Domain Suffix  : pubs.outdoors.com]
[DNS server     : 10.28.229.10]

VLAN          | IP Config | IP Address | Subnet Mask
-----+-----
DEFAULT_VLAN | Manual    | 10.28.192.1 | 255.255.255.0
```

DNS Resolver Configuration in the show ip command output

Figure C-41. Example of Viewing the Current DNS Configuration

Operating Notes

- Configuring another IP address for a priority that has already been assigned to an IP address is not allowed. To replace one IP address at a given priority level with another address having the same priority, you must first use the **no** form of the command to remove the unwanted address. Also, only one instance of a given server address is allowed in the server list. Attempting to enter a duplicate of an existing entry at a different priority level is not allowed. To change the priority of an existing server address, use the **no** form of the command to remove the entry, then re-enter the address with the new priority.
- To change the position of an address already configured with priority *x*, you must first use **no ip dns server-address priority *x* < ip-addr >** to remove the address from the configuration, then use **ip dns server-address priority < ip-addr >** to reconfigure the address with the new priority. Also, if the priority to which you want to move an address is already used in the configuration for another address, you must first use the no form of the command to remove the current address from the target priority.
- The DNS server(s) and domain configured on the switch must be accessible to the switch, but it is not necessary for any intermediate devices between the switch and the DNS server to be configured to support DNS operation.
- When multiple DNS servers are configured on the switch, they can reside in the same domain or different domains.
- A DNS configuration must include the IP address for a DNS server that is able to resolve host names for the desired domain. If a DNS server has limited knowledge of other domains, then its ability to resolve DNS-compatible command requests is also limited.
- If the DNS configuration includes a DNS server IP address but does not also include a domain suffix, then any DNS-compatible commands should include the target host's fully qualified domain name. Refer to Figure C-35 on page C-88.
- Switch-Initiated DNS packets go out through the VLAN having the best route to the DNS server, even if a Management VLAN has been configured.
- The DNS server address must be manually input. It is not automatically determined via DHCP.

Event Log Messages

Message	Meaning
DNS server address not configured	The switch does not have an IP address configured for the DNS server.
DNS server not responding	The DNS server failed to respond or is unreachable. An incorrect server IP address can produce this result.
Unknown host < <i>host-name</i> >	The host name did not resolve to an IP address. Some reasons for this occurring include: <ul style="list-style-type: none">• The host name was not found.• The named domain was not found.• The domain suffix was expected, but has not been configured. (If the server's IP address has been configured in the switch but the domain name has not been configured, then the host's fully qualified domain name must be used.)

Locator LED (Locating a Switch)

To locate where a particular switch is physically installed, use the **chassislocate** command to activate the blue Locator LED on the switch's front panel.

Syntax: chassislocate [blink | on | off]

Locates a switch by using the blue Locate LED on the front panel.

blink <1-1440>

Blinks the chassis Locate LED for a specified number of minutes (Default: 30 minutes).

on <1-1440>

Turns the chassis Locate LED on for a specified number of minutes (Default: 30 minutes).

off

Turns the chassis Locate LED off.

```
ProCurve(config)# chassislocate
  blink <1-1440>          Blink the chassis locate led (default 30 minutes).
  off                    Turn the chassis locate led off.
  on <1-1440>           Turn the chassis locate led on (default 30 minutes).
ProCurve(config)# chassislocate
```

Figure C-42. Locating a Switch with the chassislocate Command

For redundant management systems, if the active management module fails-over, the Locator LED does not remain lit.

MAC Address Management

Contents

Overview	D-2
Determining MAC Addresses	D-3
Menu: Viewing the Switch's MAC Addresses	D-4
CLI: Viewing the Port and VLAN MAC Addresses	D-5
Viewing the MAC Addresses of Connected Devices	D-7

Overview

The switch assigns MAC addresses in these areas:

- For management functions, one Base MAC address is assigned to the default VLAN (VID = 1). (All VLANs on the switches covered in this guide use the same MAC address.)
- For internal switch operations: One MAC address per port (Refer to “CLI: Viewing the Port and VLAN MAC Addresses” on page D-5.)

MAC addresses are assigned at the factory. The switch automatically implements these addresses for VLANs and ports as they are added to the switch.

Note

The switch’s base MAC address is also printed on a label affixed to the switch.

Determining MAC Addresses

MAC Address Viewing Methods

Feature	Default	Menu	CLI	Web
view switch's base (default vlan) MAC address and the addressing for any added VLANs	n/a	D-4	D-5	—
view port MAC addresses (hexadecimal format)	n/a	—	D-5	—

- **Use the menu interface** to view the switch's base MAC address and the MAC address assigned to any VLAN you have configured on the switch. (The same MAC address is assigned to VLAN1 and all other VLANs configured on the switch.)

Note

The switch's base MAC address is used for the default VLAN (VID = 1) that is always available on the switch. This is true for dynamic VLANs as well; the base MAC address is the same across all VLANs.

- **Use the CLI** to view the switch's port MAC addresses in hexadecimal format.

Menu: Viewing the Switch's MAC Addresses

The Management Address Information screen lists the MAC addresses for:

- Base switch (default VLAN; VID = 1)
- Any additional VLANs configured on the switch.

Also, the Base MAC address appears on a label on the back of the switch.

Note

The Base MAC address is used by the first (default) VLAN in the switch. This is usually the VLAN named "DEFAULT_VLAN" unless the name has been changed (by using the VLAN Names screen). On the switches covered in this guide, the VID (VLAN identification number) for the default VLAN is always "1", *and cannot be changed*.

To View the MAC Address (and IP Address) assignments for VLANs Configured on the Switch:

1. From the Main Menu, Select

- 1. Status and Counters**

- 2. Switch Management Address Information**

If the switch has only the default VLAN, the following screen appears. If the switch has multiple static VLANs, each is listed with its address data.

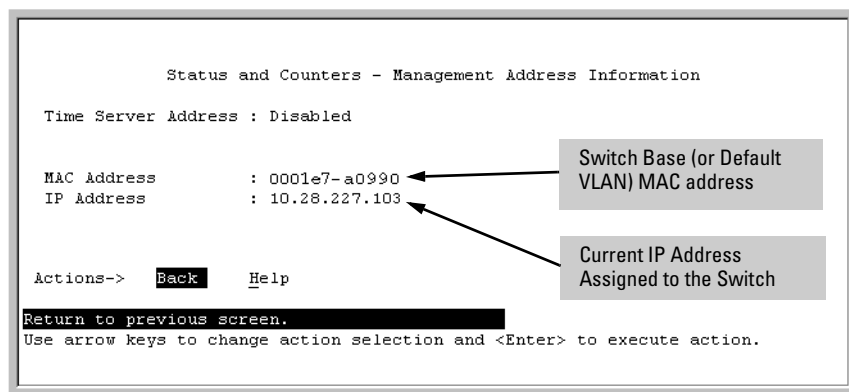


Figure D-1. Example of the Management Address Information Screen

CLI: Viewing the Port and VLAN MAC Addresses

The MAC address assigned to each switch port is used internally by such features as Flow Control and the spanning-tree protocol. Using the **walkmib** command to determine the MAC address assignments for individual ports can sometimes be useful when diagnosing switch operation.

Switch Series	MAC Address Allocation
8212zl	The switch allots 24 MAC addresses per slot. For a given slot, if a four-port module is installed, then the switch uses the first four MAC addresses in the allotment for that slot, and the remaining 18 MAC addresses are unused. If a 24-port module is installed, the switch uses the first 24 MAC addresses in the allotment, and so-on.
All Models	The switch's base MAC address is assigned to VLAN (VID) 1 and appears in the walkmib listing after the MAC addresses for the ports. (All VLANs in the switch have the same MAC address.)

To display the switch's MAC addresses, use the **walkmib** command at the command prompt:

Note

This procedure displays the MAC addresses for all ports and existing VLANs in the switch, regardless of which VLAN you select.

1. If the switch is at the CLI Operator level, use the **enable** command to enter the Manager level of the CLI.
2. Type the following command to display the MAC address for each port on the switch:

```
ProCurve# walkmib ifPhysAddress
```

(The above command is not case-sensitive.)

For example, a ProCurve 8212zl switch with the following module configuration shows MAC address assignments similar to those shown in figure D-2:

- a 4-port module in slot A, a 24-port module in slot C, and no modules in slots B and D
- two non-default VLANs configured

MAC Address Management

Determining MAC Addresses

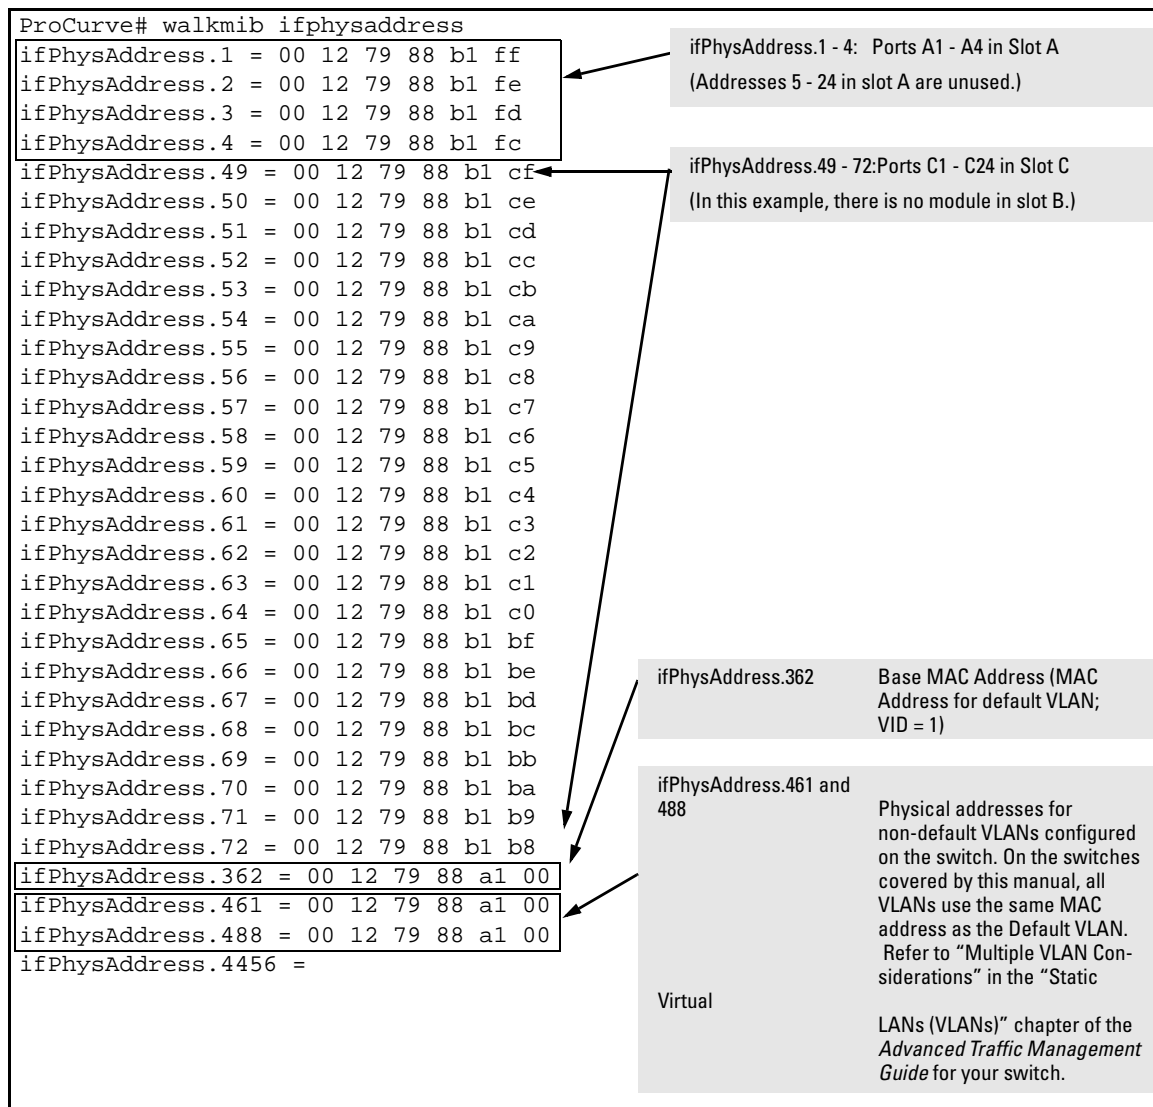


Figure D-2. Example of Port MAC Address Assignments on a Switch

Viewing the MAC Addresses of Connected Devices

Syntax: show mac-address [| *mac-addr* |

Lists the MAC addresses of the devices the switch has detected, along with the number of the specific port on which each MAC address was detected.

[*port-list*]

Lists the MAC addresses of the devices the switch has detected, on the specified port(s).

[*mac-addr*]

Lists the port on which the switch detects the specified MAC address. Returns the following message if the specified MAC address is not detected on any port in the switch:

MAC address <*mac-addr*> not found.

[vlan <*vid*>]

Lists the MAC addresses of the devices the switch has detected on ports belonging to the specified VLAN, along with the number of the specific port on which each MAC address was detected.

To list the MAC addresses of devices the switch has detected, use the **show mac-address** command.

MAC Address Management

Viewing the MAC Addresses of Connected Devices

Monitoring Resources

Contents

Viewing Information on Resource Usage	E-2
Policy Enforcement Engine	E-2
Displaying Current Resource Usage	E-4
When Insufficient Resources Are Available	E-7

Viewing Information on Resource Usage

The switch allows you to view information about the current usage and availability of resources in the Policy Enforcement engine, including the following software features:

- Access control lists (ACL)
- Quality-of-service (QoS), including device and application port priority, ICMP rate-limiting, and QoS policies
- Dynamic assignment of per-port or per-user ACLs and QoS through RADIUS authentication designated as “IDM”, with or without the optional identity-driven management (IDM) application
- Virus throttling (VT) using connection-rate filtering
- Mirroring policies, including switch configuration as an endpoint for remote intelligent mirroring
- Other features, including:
 - Management VLAN
 - DHCP snooping
 - Dynamic ARP protection
 - Jumbo IP-MTU

Policy Enforcement Engine

The Policy Enforcement engine is the hardware element in the switch that manages quality-of-service, mirroring, and ACL policies as well as other software features, using the rules that you configure. Resource usage in the Policy Enforcement engine is based on how these features are configured on the switch.

- Resource usage by dynamic port ACLs and virus-throttling is determined as follows:
 - Dynamic port ACLs configured by a RADIUS server (with or without the optional IDM application) for an authenticated client determine the current resource consumption for this feature on a specified slot. When a client session ends, the resources in use for that client become available for other uses.
 - A virus-throttling configuration (connection-rate filtering) on the switch does not affect switch resources unless traffic behavior has triggered either a throttling or blocking action on the traffic from one or more clients. When the throttling action ceases or a blocked client is unblocked, the resources used for that action are released.

- When the following features are configured globally or per-VLAN, resource usage is applied across all port groups or all slots with installed modules:
 - ACLs
 - QoS configurations that use the following commands:
 - QoS device priority (IP Address) through the CLI using the **qos device-priority** command
 - QoS application port through the CLI using **qos tcp-port** or **qos udp-port**
 - VLAN QoS Policies through the CLI using **service-policy**
 - Management VLAN configuration
 - DHCP snooping
 - Dynamic ARP protection
 - Remote mirroring endpoint configuration
 - Mirror policies per VLAN through the CLI using **monitor service**
 - Jumbo IP-MTU
- When the following features are configured per-port, resource usage is applied only to the slot or port group on which the feature is configured:
 - ACLs or QoS applied per-port or per-user through RADIUS authentication
 - ACLs applied per-port through the CLI using the **ip access-group** or **ipv6 traffic-filter** commands
 - QoS policies applied per port through the CLI using the **service-policy** command
 - Mirror policies applied per-port through the CLI using the **monitor all service** and **service-policy** commands
 - ICMP rate-limiting through the CLI using the **rate-limit icmp** command
 - Virus throttling applied to any port (when a high connection-rate client is being throttled or blocked)

Displaying Current Resource Usage

To display current resource usage in the switch, enter the **show** <qos | access-list | policy> **resources** command.

The **show resources** command output allows you to view current resource usage and, if necessary, prioritize and reconfigure software features to free resources reserved for less important features.

The **qos**, **access-list**, and **policy** parameters display the same command output and provide different ways to access task-specific information.

Syntax: show <qos | access-list | policy> resources

Displays the resource usage of the Policy Enforcement Engine on the switch by software feature. For each type of resource, the amount still available and the amount used by each software feature is shown.

Figure E1 shows the resource usage on a 3500y1 switch configured for ACLs, QoS, RADIUS-based authentication, and other features:

- The “Rules Used” columns show that ACLs, virus-throttling (VT), mirroring, and other features (for example, Management VLAN) have been configured globally or per-VLAN because identical resource consumption is displayed for each port range in the switch. If ACLs were configured per-port, the number of rules used in each port range would be different.
- The switch is also configured for virus throttling, and is either blocking or throttling routed traffic with a high rate of connection requests.
- Varying ICMP rate-limiting configurations on ports 1-24, on ports 25-48, and on slot A, have resulted in different meter usage and different rule usage listed under QoS. Global QoS settings would otherwise result in identical resource consumption on each port range in the switch.
- There is authenticated client usage of IDM resources on ports 25-48.


```

ProCurve# show qos resources

Resource usage in Policy Enforcement Engine

  Ports |      Rules      | Rules Used
  -----+-----+-----+-----+-----+-----+-----+-----+-----
  Ports | Available | ACL | QoS | IDM | VT | Mirror | Other |
  -----+-----+-----+-----+-----+-----+-----+-----+-----
  1-24 |      3014 |   15 |   11 |   0 |   1 |   0 |   3 |
  25-48 |      3005 |   15 |   10 |   10 |   1 |   0 |   3 |
  A     |      3017 |   15 |    8 |    0 |   1 |   0 |   3 |

  Ports |      Meters      | Meters Used
  -----+-----+-----+-----+-----+-----+-----+-----+-----
  Ports | Available | ACL | QoS | IDM | VT | Mirror | Other |
  -----+-----+-----+-----+-----+-----+-----+-----+-----
  1-24 |      250 |    |    5 |    0 |    |    |    0 |
  25-48 |      251 |    |    4 |    0 |    |    |    0 |
  A     |      253 |    |    2 |    0 |    |    |    0 |

  Ports | Application | Application Port Ranges Used
  -----+-----+-----+-----+-----+-----+-----+-----+-----
  Ports | Available | ACL | QoS | IDM | VT | Mirror | Other |
  -----+-----+-----+-----+-----+-----+-----+-----+-----
  1-24 |      3014 |   2 |   0 |   0 |    |    0 |   0 |
  25-48 |      3005 |   2 |   0 |   0 |    |    0 |   0 |
  A     |      3017 |   2 |   0 |   0 |    |    0 |   0 |

0 of 8 Policy Engine management resources used.
Key:
ACL = Access Control Lists
QoS = Device & Application Port Priority, QoS Policies, ICMP rate limits
IDM = Identity Driven Management
VT = Virus Throttling blocks
Mirror = Mirror Policies, Remote Intelligent Mirror endpoints
Other = Management VLAN, DHCP Snooping, ARP Protection, Jumbo IP-MTU.

Resource usage includes resources actually in use, or reserved for future
use by the listed feature. Internal dedicated-purpose resources, such as
port bandwidth limits or VLAN QoS priority, are not included.

```

Figure E1. Example of Displaying Current Resource Usage on a Series 3500yl Switch

Usage Notes for show resources Output

- A 1:1 mapping of internal rules to configured policies in the switch does not necessarily exist. As a result, displaying current resource usage is the most reliable method for keeping track of available resources. Also, because some internal resources are used by multiple features, deleting a feature configuration may not increase the amount of available resources.
 - Resource usage includes resources actually in use or reserved for future use by the listed features.
 - “Internal dedicated-purpose resources” include the following features:
 - Per-port ingress and egress rate limiting through the CLI using **rate-limit in/out**
 - Per-port ingress and egress broadcast rate limiting through the CLI using **rate-limit bcast/mcast**
 - Per-port or per-vlan priority or DSCP through the CLI using **qos priority** or **qos dscp**
 - Per protocol priority through the CLI using **qos protocol**
 - For chassis products (for example, the 5400zl or 8212zl switches), ‘slots’ are listed instead of ‘ports’ with resources shown for all installed modules on the chassis.
 - The “Available” columns display the resources available for additional feature use.
 - The “IDM” column shows the resources used for RADIUS-based authentication with or without the IDM option.
 - “Meters” are used when applying either ICMP rate-limiting or a QoS policy with a rate-limit class action.
-

When Insufficient Resources Are Available

The switch has ample resources for configuring features and supporting:

- RADIUS-authenticated clients (with or without the optional IDM application)
- Virus throttling and blocking on individual clients.

Note

Virus throttling does not operate on IPv6 traffic.

If the resources supporting these features become fully subscribed:

- The current feature configuration, RADIUS-authenticated client sessions, and virus throttling instances continue to operate normally.
- The switch generates an event log notice to say that current resources are fully subscribed.
- Currently engaged resources must be released before any of the following actions are supported:
 - Modifying currently configured ACLs, IDM, virus throttling, and other software features, such as Management VLAN, DHCP snooping, and dynamic ARP protection.

You can modify currently configured classifier-base QoS and mirroring policies if a policy has not been applied to an interface. However, sufficient resources must be available when you apply a configured policy to an interface.

- Acceptance of new RADIUS-based client authentication requests (displayed as a new resource entry for IDM).

Failure to authenticate a client that presents valid credentials may indicate that insufficient resources are available for the features configured for the client in the RADIUS server. To troubleshoot, check the event log.

- Throttling or blocking of newly detected clients with a high rate of connection requests (as defined by the current virus-throttling configuration).

The switch continues to generate event log notifications (and SNMP trap notification, if configured) for new instances of high connection-rate behavior detected by the virus-throttling feature.

Monitoring Resources
When Insufficient Resources Are Available

Daylight Savings Time on ProCurve Switches

This information applies to the following ProCurve switches:

- 212M
- 224M
- 1600M
- 2400M
- 2424M
- 4000M
- 8000M
- Series 2500
- Series 2510
- Series 2600
- Series 2610
- Series 2800
- Switch 2910
- Series 3400cl
- Series 3500
- Series 3500yl
- Series 4100gl
- Series 4200vl
- Series 5300xl
- Series 5400zl
- Switch 6108
- Switch 6200yl
- Series 6400cl
- Switch 6600
- Series 8200zl
- ProCurve AdvanceStack Switches
- ProCurve AdvanceStack Routers

ProCurve switches provide a way to automatically adjust the system clock for Daylight Savings Time (DST) changes. To use this feature you define the month and date to begin and to end the change from standard time. In addition to the value “none” (no time changes), there are five pre-defined settings, named:

- Alaska
- Canada and Continental US
- Middle Europe and Portugal
- Southern Hemisphere
- Western Europe

The pre-defined settings follow these rules:

Alaska:

- Begin DST at 2am on the second Sunday in March.
- End DST at 2am on the first Sunday in November.

Canada and Continental US:

- Begin DST at 2am on the second Sunday in March.
- End DST at 2am on the first Sunday in November.

Middle Europe and Portugal:

- Begin DST at 2am the first Sunday on or after March 25th.
- End DST at 2am the first Sunday on or after September 24th.

Southern Hemisphere:

- Begin DST at 2am the first Sunday on or after October 25th.
- End DST at 2am the first Sunday on or after March 1st.

Western Europe:

- Begin DST at 2am the first Sunday on or after March 23rd.
- End DST at 2am the first Sunday on or after October 23rd.

A sixth option named “User defined” allows you to customize the DST configuration by entering the beginning month and date plus the ending month and date for the time change. The menu interface screen looks like this (all month/date entries are at their default values):

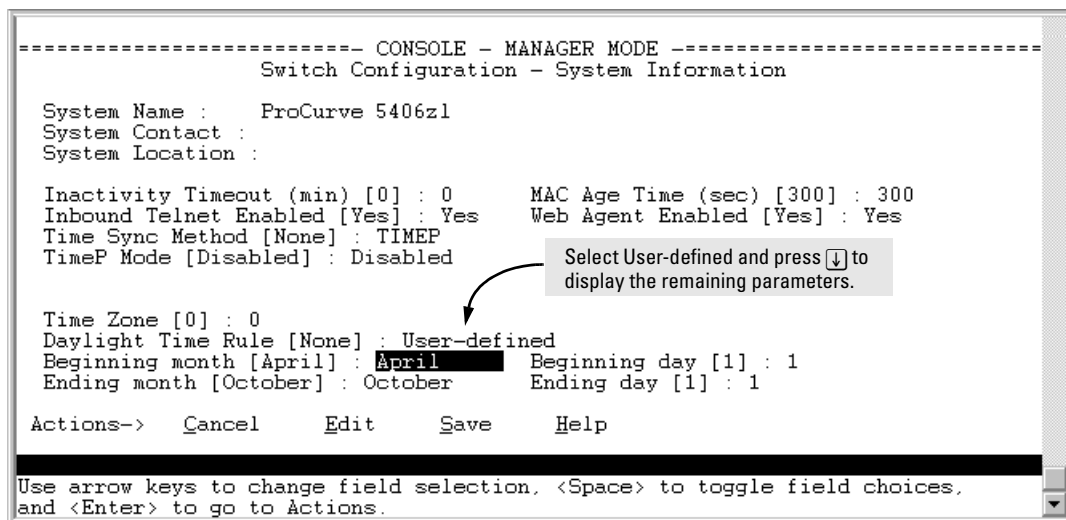


Figure F-1. Menu Interface with “User-Defined” Daylight Time Rule Option

Before configuring a “User defined” Daylight Time Rule, it is important to understand how the switch treats the entries. The switch knows which dates are Sundays, and uses an algorithm to determine on which date to change the system clock, given the configured “Beginning day” and “Ending day”:

- If the configured day is a Sunday, the time changes at 2am on that day.
- If the configured day is not a Sunday, the time changes at 2am on the first Sunday after the configured day.

This is true for both the “Beginning day” and the “Ending day”.

With that algorithm, one should use the value “1” to represent “first Sunday of the month”, and a value equal to “number of days in the month minus 6” to represent “last Sunday of the month”. This allows a single configuration for every year, no matter what date is the appropriate Sunday to change the clock.

Daylight Savings Time on ProCurve Switches

Scalability: IP Address, VLAN, and Routing Maximum Values

The following table lists the switch scalability values for the areas of VLANs, ACLs, hardware, ARP, and routing.

Subject	Maximum
IPv4 ACLs	
total named (extended or standard)	up to 2048 (minus any IPv4 numeric standard or extended ACL assignments and any RADIUS-assigned ACLs) ¹
total numeric standard	up to 99 ¹
total numeric extended	up to 100 ¹
total ACEs in all IPv4 ACLs	up to 3072 ¹
IPv6 ACLs	
total IPv6 ACLs	up to 2048 ¹
total ACEs in all IPv6 ACLs	up to 3072 ¹
¹ Actual availability depends on combined resource usage on the switch. Refer to Appendix E, "Monitoring Resources".	
Layer-3	
VLANs with at least one IP Address	512
IP addresses per system	2048 IP 2048 IPv6 ²
IP addresses per VLAN	32 ³
static routes	256
supported routes	10,000 (including ARP)
² These limits apply only to user-configured addresses and not to auto-configured link local and prefix IPv6 addresses. A maximum configuration could support up to 2048 user-configured and 2048 auto-configured IPv6 addresses for a total of 4096.	
³ There can be up to 32 IPv4 and 32 user-configured IPv6 addresses on a single VLAN. In addition, each VLAN is limited to 3 auto-configured prefix-based IPv6 addresses.	
IPv4 host hardware table	72K (8K internal, 64K external)
IPv4 BMP hardware table	2K

Scalability: IP Address, VLAN, and Routing Maximum Values

Subject	Maximum
ARP	
ARP entries	10,000
packets held for ARP resolution	25
Routing Protocol	
RIP interfaces	128
OSPF passive interfaces	512 (minus OSPF active interfaces)
OSPF active interfaces	128
OSPF areas	16
ECMP next hops	4

Switch Licensing

Switch software licensing enables advanced features in certain ProCurve switches. The following table shows the software licenses available for the switches covered by this manual.

License Type	Premium (includes OSPF, PIM – sparse mode, PIM – dense mode, VRRP, QinQ)
Switch Family	License Product
3500 and 3500yl	J8993A
5400zl	J8994A
6200yl	(included in switch)
6600	J9305A
8200zl	J9474A

General Procedure

The general procedure for installing a software license involves several different numbers:

- registration ID — This number comes with the license you purchase, and represents your right to install the particular type of license on a particular type of switch.
- hardware ID — This number is provided by the switch that you are licensing, and includes the switch's serial number and an identifier for the feature that you are licensing.
- license key — This number is generated by the My ProCurve portal, based on the registration ID and the hardware ID that you provide. When you install this number into the switch, it enables the feature that you are licensing.

The procedure for installing a licensed feature into a switch is:

1. **Locate the registration ID.** When you purchase a software license, you receive a folded license registration card. The registration ID is located on the inside of the card, in the upper left corner.
2. **Get the switch's hardware ID.** Establish a console connection to the switch CLI and enter Manager level, using the **enable** command if necessary and the switch password if required. For example:

```
ProCurve> enable
ProCurve#
```

From the Manager level, issue the **licenses hardware-id <license_type>** command. For example:

```
ProCurve# licenses hardware-id premium
```

The CLI returns a hardware ID number. Copy the hardware ID number from the screen (using Ctrl-C) or write it down. (Copying the number is easier and more accurate.) You will enter the number on the My ProCurve portal in the next step.

3. **Get the license key.** Point your Web browser at the My ProCurve portal (<http://my.procurve.com>) and sign in. Click the My Licenses tab (formerly My Software), follow the links for licensing device software, and follow the instructions for entering the registration ID and the hardware ID. At the end of the procedure a license key is displayed. (It is also e-mailed to you.) Copy the license key from the screen (using Ctrl-C) or write it down.
4. **Enter the license key into the switch.** On the CLI console, save the configuration of the switch (**write memory**). Then, from a Manager-level prompt, issue a **licenses install premium <license-key>** command. (The license key number is not case sensitive.) For example:

```
ProCurve# licenses install premium AA000GG000-A-
0123ABC-ABCD123-0A2B3C4-0123ABC
```

5. Reboot the switch. For example:

```
ProCurve# boot
or:
ProCurve# reload
```

The licensed features should now be active on the switch.

Power-Saving Features

Contents

Overview	I-2
Configuring the Power-Saving Options	I-3
Configuring the Savepower module Option	I-3
Configuring the Savepower LED Option	I-4
Configuring the Savepower port-low-pwr Option	I-6
Show Savepower Commands	I-6

Overview

There are several power-saving features that can be configured for the indicated switches and modules. The power-saving features include the ability to:

- Turn slot power on or off
- Turn LED power on or off using a timer
- Slot auto low power mode

The modules support the power-saving features as indicated in the table below.

Product Number	Description	LED Power On/Off	Slot Auto Low Power Mode	Slot Power On/Off
J8702A	ProCurve Switch zl 24 10/100/1000 PoE Module	Yes	Yes	Yes
J8705A	ProCurve Switch zl 20 Gig-T + 4 mGBIC Module	Yes	Yes	Yes
J8706A	ProCurve Switch zl 24-Port Mini-GBIC Module	Yes	No	Yes
J8707A	ProCurve Switch zl 4-Port 10GbE X2 Module	Yes	No	Yes
J8708A	ProCurve Switch zl 4-Port 10GbE CX4 Module	Yes	No	Yes
J9307A	HP ProCurve 24-Port 10/100/1000 PoE+ zl Module	Yes	Yes	Yes
J9308A	HP ProCurve 20-Port 10/100/1000 PoE+/4-Port MiniGBIC zl Module	Yes	Yes	Yes
J9309A	HP ProCurve 4-Port 10Gbe SFP+ zl Module	Yes	No	Yes
J9478A	HP ProCurve 24-Port 10/100 PoE+ zl Module	Yes	Yes	Yes

Configuring the Power-Saving Options

The **savepower** command provides configurable power-saving options.

Syntax: [no] savepower <module [slot-list | all] | led [slot-id] | port-low-pwr [slot-id]>

Configures power-saving features.

module [slot-id]: Turns power-saving options on or off for all modules or a specified module.

*The **no** form of the command powers on all the slots if they are powered off already.*

led [slot-id]: Turns power-saving options on or off for the LEDs for all modules or a specified module.

port-low-pwr [slot-id]: Enables or disables auto power down for all slots or a specified slot.

Configuring the Savepower module Option

The **module** option provides the ability to turn the slot power on or off. If no module is specified, then all slots are powered off. You can also specify **all** to turn off the power for all slots. If the command is preceded by **no**, then all the slots are powered on, if off already.

```
ProCurve(config)# savepower module c
ProCurve(config)# show savepower module

Module Save Power Information

Slot | Status
----+-----
  A  | Disabled
  B  | Disabled
  C  | Enabled
  D  | Disabled
  E  | Disabled
```

Figure I-1. Example of savepower module Command

The **savepower module** command shuts down the specified modules in the order specified in the command. The ports on these modules no longer pass traffic. Any management traffic (SNMP, SSH, Telnet) that passes through these modules is interrupted. It can take up to two minutes to power down all the specified modules. Check the event log to see the current status of the module power down. This command applies to PoE/PoE+ modules as well as non-PoE/PoE+ modules.

You can verify the status of the **savepower** command by using the **show modules** command or by checking the log messages (for 8200zl and 5400zl switches).

Note

If a **savepower module <slot-list>** or **savepower all** command is immediately followed by a **no savepower module <slot-list>** or **no savepower all** command, the first slot in the list is powered down and then brought up.

Configuring the Savepower LED Option

The savepower LED option provides the ability to configure a timer for turning off the chassis LEDs as well as the configured slot LEDs. There is one system-wide timer; all the selected slots will have the chassis LEDs turned off for the same amount of time.

Syntax: [no] savepower led [slot-id] < MM/DD/[YY]YY <HH:MM> | now > duration [HH:<MM> [recur]

Schedules a timer for turning off the chassis LEDs and configured slot LEDs. The LEDs are turned off for the configured time period and duration.

If a slot is specified, the LEDs for that slot are turned off. This is enabled by the timer command, however, if a timer is already running, the feature is enabled immediately.

*The **all** option can be specified for the **slot-id**. All the switch LEDs are turned off.*

<MM/DD/[YY]YY <HH:MM>>: Specifies the date and time to start the timer.

now: Instantaneously turns off the LEDs. The configured timer is canceled and all the configured modules go into power-saving mode immediately.

duration <[HH:]MM>: The amount of time the LEDs remain turned off. Optional. If the duration value is zero, when the timer starts the LEDs are turned off indefinitely until the timer is canceled or the command is overridden with another command. Default: 0 (zero)

*recur: Optional. If specified, the LEDs are turned off on a daily basis at the configured time. The **recur** option is ignored if the duration is configured as zero. Default: disabled.*

If the configured time is less than two minutes from the current time, the LEDs will be turned off instantly, however, the start time of the timer is shown as two minutes from the current time.

A new command overrides the previous command, regardless of the current state. For example, if a timer is active and new command is given, the currently running timer is canceled and the new timer is scheduled.

The **no** form of the **savepower led** command cancels any scheduled or running timer and the LEDs are returned to their original state. The **all** option can be specified with the **no** command to turn on all the switch LEDs.

```
ProCurve(config)# savepower led timer 06/01/2009 12:01 duration 12:00 recur

ProCurve(config)# show savepower led

Led Save Power Information

Alarm Start Time       : 06/01/09 12:01:07
Alarm Duration (HH:MM) : 12:00
Recurrent Status      : Enabled

Led Save Power Information

Slot | Status
----+-----
A   | Disabled
B   | Disabled
C   | Disabled
D   | Disabled
E   | Disabled
```

Figure I-2. Example of Setting a Time and Duration for savepower led Command

Configuring the Savepower port-low-pwr Option

The **port-low-pwr** option puts the slots into auto low power mode if they are not linked. If a particular slot is specified, only that slot goes into auto low power mode. Specifying **all** puts all the slots into auto low power mode.

The ports in low power mode periodically monitor to determine if the link has become active. If a LAN cable is connected to one of the ports, that port will come out of the low power mode state after approximately 2 seconds (the monitor period) and enter into normal power mode. The remaining ports continue to be in low power mode.

The **no** form of the command puts the specified slot into normal power mode. Specifying **all** with the **no** form of the command puts all the slots into normal power mode.

```
ProCurve(config)# savepower port-low-pwr c
ProCurve(config)# show savepower port-low-pwr

Port Save Power Information

Slot | Status
----+-----
A    | Disabled
B    | Disabled
C    | Enabled
D    | Disabled
E    | Disabled
```

Figure I-3. Example of savepower port-low-power Command for Slot C

Show Savepower Commands

The settings for the **savepower** commands can be viewed using the appropriate **show** command.

Show Savepower Module.

To display the settings for the **savepower module** command, use **show savepower module**.

```

ProCurve(config)# show savepower module

Module Save Power Information

Slot | Status
-----+-----
A    | Disabled
B    | Disabled
C    | Enabled
D    | Disabled
E    | Disabled

```

Figure I-4. Example of Output for show savepower module Command

Show Savepower Port-low-pwr.

To display the status of the power-down feature for the slots, use the **show savepower port-low-pwr** command. For the stackable switches, the output shows if the feature is enabled or not enabled.

```

ProCurve(config)# show savepower port-low-pwr

Port Save Power Information

Slot | Status
-----+-----
A    | Enabled
B    | Enabled
C    | Enabled
D    | Enabled
E    | Enabled

```

Figure I-5. Example of Output for show savepower port-low-pwr Command

Show Savepower LED.

To display the configured status of the LED power-saving option, use the **show savepower led** command.

```
ProCurve(config)# show savepower led

Led Save Power Information

Alarm Start Time      : 06/01/09 12:01:07
Alarm Duration (HH:MM) : 12:00
Recurrent Status      : Enabled

Led Save Power Information

Slot | Status
----+-----
A    | Enabled
B    | Enabled
C    | Enabled
D    | Enabled
E    | Enabled
```

Figure I-6. Example of Output for show savepower led Command

Network Out-of-Band Management (OOBM) for the 6600 Switch

Contents

Concepts	J-2
Example	J-4
OOBM and Switch Applications	J-5
Tasks	J-6
OOBM Configuration	J-6
OOBM Context	J-6
OOBM Enable/disable	J-7
OOBM Port Enable/disable	J-8
OOBM Port Speed Control	J-9
OOBM IPv4 Address Configuration	J-10
OOBM IPv4 Default Gateway Configuration	J-10
OOBM Show Commands	J-11
Show OOBM	J-11
Show OOBM IP Configuration	J-12
Show OOBM ARP Information	J-12
Application Server Commands	J-13
Application Client Commands	J-14
Example	J-15

Concepts

Management communications with a managed switch can be:

- in band—through the networked data ports of the switch
- out of band—through a dedicated management port (or ports) separate from the data ports

Out-of-band ports have typically been serial console ports using DB-9 or specially wired 8-pin modular (RJ-style) connectors. Some recent HP ProCurve switches have added networked out-of-band management ports. The illustration below shows management connections for a typical switch.




Figure J-1. Management ports

Out-of-band management (OOBM) operates on a “management plane” that is separate from the “data plane” used by data traffic on the switch and by in-band management traffic. That separation means that out-of-band management can continue to function even during periods of traffic congestion, equipment malfunction, or attacks on the network. In addition, it can provide

improved switch security: a properly configured switch can limit management access to the management port only, preventing malicious attempts to gain access via the data ports.

Network OOBM typically occurs on a management network that connects multiple switches. It has the added advantage that it can be done from a central location and does not require an individual physical cable from the management station to each switch's console port.

Of the switches covered by this manual, network OOBM is available on:

- HP ProCurve 6600-24XG switch (J9265A)
- HP ProCurve 6600-48G switch (J9451A)
- HP ProCurve 6600-48G-4XG switch (J9452A)

The table below summarizes the switch management ports.

Table J-1. Switch Management Ports

	In Band	Out Of Band	
	Networked	Directly connected	Networked
Management interface	command line (CLI), menu, Web	command line (CLI), menu	command line (CLI), menu
Communication plane	data plane	management plane	management plane
Connection port	any data port	dedicated serial or USB console port	dedicated networked management port
Connector type	usually RJ-45; also CX4, SFP, SFP+, and XFP	DB9 serial, serial-wired 8-pin RJ	RJ-45
Advantages	allows centralized management	not affected by events on data network, shows boot sequence	not affected by events on data network; allows centralized management; allows improved security
Disadvantages	can be affected by events on data network; does not show boot sequence	requires direct connection to console port (can be done via networked terminal server)	does not show boot sequence

Example

In a typical data center installation, top-of-rack switches connect servers to the data network, while the management ports of those switches connect to a physically and logically separate management network. This allows network administrators to manage the switches even if operation on the data network is disrupted.

In the illustration below, the switches face the hot aisle of the data center, allowing easy connection to the network ports on the backs of the servers.




Figure J-2. Network out-of-band management in a data center

For even more control, the serial console ports of the switches could be connected to the management network through a serial console server (essentially, a networked serial switch), allowing the network administrators to view the CLI activity of each switch at boot time and to control the switches through the console ports (as well as through the management ports).

OOBM and Switch Applications

The table below shows the switch applications that are supported on the OOBM interface as well as on the data interfaces. In this list, some applications are client-only, some are server-only, and some are both.

Application	Inbound OOBM (server)	Outbound OOBM (client)	Inbound Data Plane (server)	Outbound Data Plane (client)
Telnet	yes	yes	yes	yes
SSH	yes	n/a	yes	n/a
SNMP	yes	yes*	yes	yes
TFTP	yes	yes	yes	yes
HTTP	yes	n/a	yes	n/a
SNTP	n/a	yes	n/a	yes
TIMEP	n/a	yes	n/a	yes
RADIUS	n/a	yes	n/a	yes
TACACS	n/a	yes	n/a	yes
DNS**	n/a	yes	n/a	yes
Syslog	n/a	yes	n/a	yes
Ping	yes***	yes	yes***	yes
Traceroute	yes***	yes	yes***	yes
n/a = not applicable * = SNMP client refers to SNMP traps as they originate from the switch. ** = DNS has a limit of two servers — primary and secondary. Either can be configured to use the OOBM interface. *** = Ping and Traceroute do not have explicit servers. Ping and Traceroute responses are sent by the host stack.				

For applications that have servers, **oobm/data/both** options have been added to listen mode. There is now a **listen** keyword in the CLI commands to allow selection of those options. Default value is **both** for all servers. See the Application Server Commands (page J-13) in the “Tasks” section below for details of the new command syntax.

Tasks

OOBM Configuration

OOBM Context

OOBM configuration commands can be issued from the global configuration context (`config`) or from a specific OOBM configuration context (`oobm`). To enter the OOBM configuration context from the general configuration context, use the **oobm** command.

Syntax: `oobm`

Enters the OOBM context from the general configuration context.

For example:

```
ProCurve (config)# oobm
ProCurve (oobm)#
```

OOBM Enable/disable

To enable or disable network OOBM, use the **enable** or **disable** command. Network OOBM is enabled by default.

Syntax:

From the OOBM context:

```
enable  
disable
```

From the general configuration context:

```
oobm enable  
oobm disable
```

Enables or disables networked out-of-band-management on the switch.

OOBM is not compatible with either a management VLAN or stacking. If you attempt to enable OOBM when a management VLAN is enabled or when stacking is enabled, the command will be rejected and you will receive an error message.

If an OOBM IP address exists and you disable OOBM, the OOBM IP address configuration is maintained. If you enable OOBM and there is a pre-existing OOBM IP address, it will be reinstated.

Examples:

```
ProCurve (oobm)# enable  
ProCurve (oobm)# disable  
ProCurve (config)# oobm enable  
ProCurve (config)# oobm disable
```

OOBM Port Enable/disable

The OOBM **interface** command enables or disables the OOBM interface (the OOBM port, as opposed to the OOBM function).

Syntax:

From the OOBM context:

```
interface [enable | disable]
```

From the general configuration context:

```
oobm interface [enable | disable]
```

Enables or disables the networked OOBM interface (port).

For example:

```
ProCurve (oobm)# interface enable
```

```
ProCurve (config)# oobm interface disable
```

OOBM Port Speed Control

The OOBM port operates at 10 Mbps or 100 Mbps, half or full duplex. These can be set explicitly or they can be automatically negotiated using the **auto** setting. Set the port speed using the **interface** command.

Syntax:

From the OOBM context:

```
interface speed-duplex [10-half | 10-full | 100-half | 100-full | auto]
```

From the general configuration context:

```
oobm interface speed-duplex [10-half | 10-full | 100-half | 100-full | auto]
```

Enables or disables the networked OOBM interface (port).

Available settings are:

10-half	10 Mbps, half-duplex
10-full	10-Mbps, full-duplex
100-half	100-Mbps, half-duplex
100-full	100-Mbps, full-duplex
auto	auto negotiate for speed and duplex

For example:

```
ProCurve (oobm)# interface speed-duplex auto
```

OOBM IPv4 Address Configuration

Configuring an IPv4 address for the OOBM interface is similar to VLAN IP address configuration, but it is accomplished within the OOBM context.

Syntax:

From the OOBM context:

```
[no] ip address [dhcp-bootp | ip-address/mask-length]
```

From the general configuration context:

```
[no] oobm ip address [dhcp-bootp | ip-address/mask-length]
```

Configures an IPv4 address for the switch's OOBM interface.

You can configure an IPv4 address even when global OOBM is disabled; that address will become effective when OOBM is enabled.

For example:

```
ProCurve (oobm)# ip address 10.1.1.17/24
```

OOBM IPv4 Default Gateway Configuration

Configuring an IPv4 default gateway for the OOBM interface is similar to VLAN default gateway configuration, but it is accomplished within the OOBM context.

Syntax:

From the OOBM context:

```
[no] ip default-gateway ip-address
```

From the general configuration context:

```
[no] oobm ip default-gateway ip-address
```

Configures an IPv4 default gateway for the switch's OOBM interface.

For example:

```
ProCurve (oobm)# ip default-gateway 10.1.1.1
```

OOBM Show Commands

The **show** commands for OOBM are similar to the analogous commands for the data plane. Note that you must always include the **oobm** parameter to see the information for the OOBM interface, regardless of the context. For instance, even from the OOBM context the **show ip** command displays the IP configuration for the data plane; to see the IP configuration of the OOBM interface you need to use **show oobm ip**.

Show OOBM

This command shows the global OOBM and OOBM port configurations.

Syntax: show oobm

Summarizes OOBM configuration information. This command displays the global OOBM configuration (enabled or disabled), the OOBM interface status (up or down) and the port status (enabled/disabled, duplex, and speed).

You can issue this command from any context.

For example:

```
ProCurve# show oobm
```

```
Global Configuration
OOBM Enabled           : Yes
OOBM Port Type         : 10/100TX
OOBM Interface Status  : Up
OOBM Port              : Enabled
OOBM Port Speed        : Auto
```

Show OOBM IP Configuration

Use **show oobm ip** to see the IP configuration of the OOBM interface.

Syntax: show oobm ip

Summarizes the IP configuration of the OOBM interface. This command displays the status of IPv4 (enabled/disabled), the IPv4 default gateway, and the IPv4 address configured for the interface.

You can issue this command from any context.

For example:

```
ProCurve# show oobm ip
```

Show OOBM ARP Information

Use **show oobm arp** to see the ARP table entries for the OOBM interface.

Syntax: show oobm arp

Summarizes the ARP table entries for the OOBM interface.

You can issue this command from any context.

```
ProCurve# show oobm arp
```


Application Server Commands

Application servers (as described in OOBM and Server Applications in the Concepts section above) have added a **listen** keyword with **oobm|data|both** options to specify which interface(s) is(are) active.

Default value is **both** for all servers.

For example:

Telnet: **telnet-server [listen <oobm | data | both>]**
Management and Configuration Guide, page 7-6

SSH: **ip ssh [listen <oobm | data | both>]**
Access Security Guide, page 8-18

SNMP: **snmp-server [listen <oobm | data | both>]**
Management and Configuration Guide, page 14-33

TFTP: **tftp server [listen <oobm | data | both>]**
Management and Configuration Guide, page A-10

HTTP: **web-management [listen <oobm | data | both>]**
Management and Configuration Guide, page 7-8

In all cases, **show running-config** will display the server configurations.

Use the **no** form of the command to prevent the server from running on either interface. For example:

Telnet: **no telnet-server**

SSH: **no ip ssh ...**

SNMP: **no snmp-server ...**

TFTP: **no tftp server**

HTTP: **no web-management ...**

The **show servers** command shows the listen mode of the servers.

```
ProCurve# show servers
```

```
Server listen mode
```

Server	Listen mode
Telnet	both
Ssh	both
Tftp	both
Web-management	both
Snmp	both

Application Client Commands

CLI commands for client applications have added the **oobm** keyword to allow you to specify that the outgoing request be issued from the OOBM interface. If you do not specify the **oobm** keyword, the request will be issued from the appropriate in-band data interface. Command syntax is:

Telnet: **telnet <ip-address> [oobm]**

Management and Configuration Guide, page 7-7

TFTP: **copy tftp ... <ip-address> <filename> ... [oobm]**

Management and Configuration Guide, page A-4 and following

SNTP: **[no] sntp server priority <priority> <ip-address> [oobm] [version]**

Management and Configuration Guide, page 9-13

TIMEP: **[no] ip timep <dhcp | manual <ip-address> [oobm] > [...]**

Management and Configuration Guide, page 9-32

RADIUS: **[no] radius-server host <ip-address> [oobm]**

Access Security Guide, page 4-18, page 6-15, page 13-26

TACACS+: **[no] tacacs-server host <ip-address> [oobm]**

Access Security Guide, page 5-19

DNS: **[no] ip dns server-address priority <priority> <ip-address> [oobm]**

Management and Configuration Guide, page C-91

Syslog: **[no] logging <ip-address> [control-descr] | [oobm]**

Management and Configuration Guide, page C-54

Ping: `ping [...] [source <ip-address | vlan-id | oobm>]`
Management and Configuration Guide, page C-65

Traceroute: `traceroute [...] [source <ip-address | vlan-id | oobm>]`
Management and Configuration Guide, page C-67

Example

This example shows setup and use of network OOBM using the commands described above.

Assume that the figure below describes how you want to set up your data center.




Figure J-3. Example data center

Assume that you are configuring the switch in the left-hand rack to communicate on both the data and management networks. You might do the following:

- Configure an IP address on the data network.
- Verify that out-of-band management is enabled. (It is enabled by default.)
- Configure an IP address on the management network.
- Verify that the switch can communicate on both networks.

Network Out-of-Band Management (OOBM) for the 6600 Switch Tasks

The CLI commands that follow would accomplish those tasks. (The first time through the process you might easily make the omission shown near the end of the example.)

```
Switch 41# config
Switch 41(config)# vlan 1
Switch 41(vlan-1)# ip address 10.1.129.7/20
Switch 41(vlan-1)# end
Switch 41# show oobm
```

Set up IP address on data network.
Exit back to manager context.
Look at default OOBM configuration.

```
Global Configuration
OOBM Enabled           : Yes
OOBM Port Type         : 10/100TX
OOBM Interface Status : Up
OOBM Port              : Enabled
OOBM Port Speed        : Auto
```

Defaults look appropriate.

```
Switch 41# config
Switch 41(config)# oobm
Switch 41(oobm)# ip address 10.255.255.41/24
Switch 41(oobm)# ip default-gateway 10.255.255.1
Switch 41(oobm)# end
Switch 41# ping 10.1.131.44
10.1.131.44 is alive, time = 19 ms
Switch 41# ping 10.1.131.51
10.1.131.51 is alive, time = 15 ms
Switch 41# ping 10.255.255.42
The destination address is unreachable.
Switch 41# ping source oobm 10.255.255.42
10.255.255.42 is alive, time = 2 ms
Switch 41#
```

Go to OOBM context and
add IP address and
default gateway.
Exit back to manager context.
Ping server in this rack (on data network).
Ping server in adjacent rack.
Ping switch in adjacent rack.
Oops! It's on the management network.
Go through the management port
and it works fine.

Index

Symbols

=> **prompt** ... C-84

Numerics

802.1X

effect, LLDP ... 14-79
LLDP blocked ... 14-46

802.1X access control

authentication failure, SNMP
notification ... 14-26
SNMP notification of authentication
failure ... 14-26

A

access

manager ... 14-13
operator ... 14-13
out-of-band ... 2-3

access control list

See ACL.

ACL

debug messages ... C-42
See also debug command.
dynamic port ACL ... E-2
gateway fails ... C-12
mirroring
 converted to classifier-based
 policies ... B-30
 replaced by classifier-based match
 criteria ... B-29, B-67
resource usage ... E-2
resources ... E-4
transferring command files ... A-36
troubleshooting ... C-9

ACL, IPv4

limit ... G-1
RADIUS-assigned, limit ... G-1
scalability ... G-1

ACL, IPv6

limit ... G-1
RADIUS-assigned, limit ... G-1

scalability ... G-1

ACLs

See ACL.

address

network manager ... 14-4

address table, port ... B-19

address, network manager ... 14-5

advertise location ... 14-58

AES encryption ... 14-9

alert log ... 5-21

alert types ... 5-22
disabling ... 5-25
setting the sensitivity level ... 5-24
sorting the entries ... 5-21

alias

command ... 4-18

allocation, class ... 11-12

allocation, value ... 11-12

ARP

arp age, default ... 8-7
maximums ... G-2

ARP protection

SNMP notification ... 14-17, 14-26

arp-protect

debug messages ... C-42

asterisk

meaning in show config ... 6-30
meaning in traceroute ... C-69

authentication

notification messages ... 14-17, 14-26
SNTP ... 9-20
SNTP client ... 9-16

authentication trap

See also SNMP.

authorized IP managers

SNMP, blocking ... 14-3

auto MDI/MDI-X

configuration, display ... 10-23
operation ... 10-21, 10-23
port mode, display ... 10-23

Auto-10 ... 12-5, 12-8, 12-19

autonegotiate ... 14-58

auto-TFTP ... A-11

disable ... A-11, A-14

- disabled ... A-11
- download to a redundant management system ... A-9
- downloading software images ... A-11

B

bandwidth

- displaying port utilization ... 10-13
- displaying utilization ... 5-18
- guaranteed minimum
 - See* guaranteed minimum bandwidth.

banner

- configuring ... 2-11
- default ... 2-9
- non-default ... 2-10
- operation ... 2-9
- redundant management ... 2-10

Best Offer ... 6-43

blue locator LED ... C-96

boot

- See also* reboot.

boot active ... 15-13

boot command ... 6-4, 6-20

boot ROM console ... A-5

boot ROM mode ... C-84

boot-history command ... 15-42

Bootp

- Bootp table file ... 8-14
- Bootptab file ... 8-14
- effect of no reply ... C-8
- operation ... 8-12, 8-14
- server ... 8-3
- using with Unix systems ... 8-14
- See also* DHCP.

Bootp/DHCP differences ... 8-13

Bootp/DHCP, LLDP ... 14-54

broadcast limit ... 10-5, 10-20

broadcast mode

- SNTP ... 9-21

broadcast storm ... 12-4, C-19

broadcast traffic

- IPX ... 10-5, 10-21
- RIP ... 10-5, 10-21

browser interface

- See* web browser interface.

C

CDP ... 14-80, 14-81, 14-82, 14-83, 14-85

chassislocate ... B-8

chassislocate LED ... C-96

Classifier

- ACL match criteria deprecated ... B-29, B-67
- benefits ... B-29, B-66
- class configuration
 - mirroring ... B-68
- IPv6 match criteria ... B-66
- mirroring configuration ... B-29, B-31, B-34, B-35, B-56, B-66, B-67
 - default class ... B-71
 - examples ... B-88
 - restrictions ... B-67, B-72, B-73
 - viewing ... B-73
- policy configuration
 - mirroring ... B-69
- resource usage, displaying ... B-84
- restrictions
 - mirroring configuration ... B-67, B-72

Clear + Reset button combination ... 6-37

Clear button ... 5-11

- restoring factory default configuration ... C-84

clear logging ... C-37

clear statistics

- global ... B-18
- ports ... B-18

CLI

- accessing from menu console ... 3-8
- context configuration level ... 4-5
- context level ... 10-16
- global configuration level ... 4-5
- Help ... 4-11
- keystroke shortcuts ... 4-20
- listing command options ... 4-8
- moving to or from the menu ... 4-7
- port or trunk-specific commands ... 4-13
- privilege levels ... 4-3
- using ... 4-2–4-16, ??–4-20
- VLAN-specific commands ... 4-15

command line interface

- See* CLI.

command syntax conventions ... 1-2

communities, SNMP ... 14-14

- viewing and configuring with the CLI ... 14-15
- viewing and configuring with the menu ... 14-13

config

- copy tftp oobm ... A-31
- config files**
 - oobm ... 6-38
- config files, SCP/SFTP transfer ... 6-40**
- configuration**
 - Bootp ... 8-14
 - clearing module ... 10-31
 - comparing startup to running ... 6-6
 - console ... 7-3
 - copying ... A-29
 - DHCP Option 66 ... 6-41
 - DHCP, Best Offer ... 6-43
 - factory default ... 6-9, 8-2
 - file update with Option 66 ... 6-41
 - file updating with Option 66 ... 6-42
 - impacts of software download on ... A-5
 - IP ... 8-2
 - Option 67 ... 6-42
 - permanent ... 6-7
 - permanent change defined ... 6-4
 - port ... 10-1
 - port trunk group ... 12-1
 - port, duplex ... 10-15
 - port, speed ... 10-15
 - quick ... 3-8
 - reboot to activate changes ... 3-13
 - restoring factory defaults ... C-83
 - saving from menu interface ... 3-10
 - serial link ... 7-3
 - SNMP ... 14-4, 14-5, 14-11
 - SNMP communities ... 14-13, 14-15
 - startup ... 3-10
 - system ... 7-12
 - Telnet access configuration ... 7-3
 - TFTP server ... 6-41
 - traffic mirroring ... B-28
 - transferring ... A-29
 - trap receivers ... 14-19
 - updating the file using Option 66 ... 6-42
 - usb autorun ... A-47
 - using Menu interface ... 3-7
 - viewing ... 6-6
 - web browser access ... 7-3
- configuration file**
 - browsing for troubleshooting ... C-71
- configuration file, multiple**
 - after first reboot ... 6-29
 - applications ... 6-27
 - asterisk ... 6-30
 - backupConfig ... 6-28
 - change policy ... 6-31
 - Clear + Reset button combination ... 6-37
 - copy from a USB device ... A-35
 - copy from tftp host ... 6-39
 - copy to a USB device ... A-35
 - copy to tftp host ... 6-38
 - copy via tftp ... A-31
 - copy via Xmodem ... A-33
 - create new file ... 6-28, 6-34, 6-35
 - current file in use ... 6-30
 - default reboot from primary ... 6-32
 - erasing ... 6-35
 - memory assignments ... 6-29
 - memory slot ... 6-27, 6-30, 6-32
 - minconfig ... 6-32, 6-36
 - newconfig ... 6-32
 - oldConfig ... 6-29
 - override reboot policy ... 6-31
 - policy, override ... 6-33
 - power cycle ... 6-32
 - primary boot path ... 6-30
 - reboot policy options ... 6-27
 - reboot policy, override ... 6-31
 - reboot process ... 6-28
 - reload ... 6-33
 - rename config file ... 6-34
 - reset ... 6-32
 - running-config file ... 6-28
 - running-config file operation ... 6-28
 - secondary boot path ... 6-30
 - show config file content ... 6-31
 - show multiple files ... 6-30
 - startup-config ... 6-28
 - startup-config file ... 6-28
 - transition to multiple files ... 6-29
 - unable to copy ... 6-34
 - workingConfig ... 6-28, 6-29
 - xmodem from host ... 6-40
 - xmodem to host ... 6-40
- connection-rate filtering**
 - affect on switch resources ... E-2
 - resource usage ... E-2
- console**
 - Actions line ... 3-10, 3-11
 - configuring ... 7-3
 - ending a session ... 3-5

- features ... 2-3
- Help ... 3-9, 3-11
- inactivity-timer ... 7-9
- Main Menu interface ... 3-7
- meaning of asterisk ... 3-10, 3-13
- measuring network activity ... C-8
- navigation ... 3-9, 3-10
- operation ... 3-10
- starting a session ... 3-4
- statistics, clear counters ... 3-12
- status and counters access ... 3-7
- status and counters menu ... B-6
- troubleshooting access problems ... C-6
- context level**
 - global config ... 4-5, 8-10
 - manager level ... 4-5
 - moving between contexts ... 4-7
 - port or trunk-group ... 4-13
 - VLAN-specific ... 4-15
- copy**
 - command output ... A-40
 - crash data ... A-41
 - crash log ... A-43
 - event log output ... A-41
 - multiple config file, tftp ... 6-38
 - software images ... A-28
 - tftp show-tech ... A-31
- copy config**
 - oobm ... A-30
- copy show tech** ... C-75
- copy tftp**
 - show-tech ... A-31
- CPU utilization** ... B-7
- cpu utilization data** ... B-9
- crash log**
 - oobm ... A-43
- custom, show tech** ... A-32
- customizing, show command output** ... 10-10

D

- date format, events** ... C-28
- date, configure** ... 7-18
- debug**
 - acl messages ... C-42
 - compared to event log ... C-41
 - destination, logging ... C-42
 - displaying debug configuration ... C-46

- forwarding IPv4 messages ... C-43
- lldp messages ... C-43
- overview
- packet messages ... C-43
- sending event log messages ... C-41
- standard event log messages ... C-43
- using CLI session ... C-42
- wireless services ... C-43

debug command

- arp-protect ... C-42
- configuring debug/Syslog operation ... C-44
- destinations ... C-42, C-52
- dhcp-snooping ... C-42
- event log ... C-61
- event log as default ... C-42
- event log messages ... C-50
- event types supported ... C-41
- operating notes ... C-60
- OSPF messages ... C-51
- RIP messages ... C-51
- show debug ... C-46
- support for "debug" severity on Syslog
 - servers ... C-52, C-61
- syntax ... C-42, C-50
- using CLI session ... C-52
- vrrp ... C-43

debug logging

- LLDP ... 14-44

default gateway ... 8-3

See also gateway.

default settings

- auto-TFTP, *disabled* ... A-11
- banner ... 2-9, 2-13
- baud rate, *speed sense* ... 7-3
- boot flash, *primary* ... 6-19
- configuration file name, *switch.cfg* ... 6-42
- console/serial configuration ... 7-5
- default gateway, *none* ... 8-2
- DHCP Option 66, *enabled* ... 6-41
- flow control, *XON/XOFF* ... 7-3
- flow-control, *disabled* ... 10-19
- inactivity timer, *0 minutes* ... 7-3
- interface access features ... 7-3
- IP address, *DHCP/Bootp* ... 8-3
- IP configuration features ... 8-2
- loopback interface, *le0* ... 8-16
- MAC age time, *300 seconds* ... 7-12
- mdix-mode, *auto-mdix* ... 10-23

- PoE ... 11-8
- PoE allocation, *usage* ... 11-12
- PoE power threshold, *80* ... 11-17
- PoE pre-std-detect, *enabled* ... 11-8
- PoE priority, *low* ... 11-10
- PoE value, *17W* ... 11-21
- port speed, *auto* ... 10-16
- security ... A-48
- SNTP ... 9-5
- sntp poll interval, *720 seconds* ... 9-11
- Support/Mgmt URL window ... 5-13
- system information features ... 7-12
- system name, *switch product name* ... 7-12
- Telnet access, *enabled* ... 7-3
- terminal type, *VT-100* ... 7-3
- TFTP, *enabled* ... A-10
- time sync method, *none* ... 7-12
- time synchronization protocol, *TimeP* ... 9-4
- time zone, *0* ... 7-13
- Time-to-Live (TTL), *64* ... 8-3
- UDLD, *disabled* ... 10-34
- usb autorun, *disabled* (if password) ... A-50
- usb autorun, *enabled* (if no password) ... A-50
- Web browser access, *enabled* ... 7-3
- default trunk type** ... 12-11
- default VLAN** ... 8-4
- DES encryption** ... 14-9
- Device Passwords window** ... 5-9
- DHCP**
 - address problems ... C-8
 - Best Offer ... 6-43
 - Bootp operation ... 8-12
 - effect of no reply ... C-8
 - manual gateway precedence ... 8-13
 - Option 66 ... 6-41
- DHCP snooping**
 - resource usage ... E-2
 - SNMP notification ... 14-17, 14-26
- DHCP/Bootp differences** ... 8-13
- DHCP/Bootp process** ... 8-13
- DHCP/Bootp, LLDP** ... 14-54
- dhcp-snooping**
 - debug messages ... C-42
- DHCPv6**
 - debug messages ... C-43
- dhcpv6-client** ... C-43
- diagnostics tools** ... C-62
 - browsing the configuration file ... C-71

- displaying switch operation ... C-72, C-75
- ping and link tests ... C-63
- traceroute ... C-67
- viewing switch operation ... C-71

distributed trunking ... 12-27

DNS

- configuration ... C-89, C-92
- configuration error ... C-95
- configuration, viewing ... C-93
- DNS-compatible commands ... C-87, C-89
- domain name, fully qualified ... C-87, C-88, C-93
- domain suffix ... C-87
- domain-name configuration ... C-91
- event log messages ... C-95
- example ... C-91
- host name ... C-87
- IPv6 DNS resolution ... C-87
- name, using in web browser
- operating notes ... C-94
- ping ... C-87, C-89, C-92
- resolver ... C-87
- resolver operation ... C-88
- secure management VLAN ... C-94
- server address, DHCP not used ... C-94
- server IP address ... C-88, C-94
- server-address configuration ... C-90
- three entries supported ... C-90
- three server entries supported ... C-90
- traceroute ... C-87, C-89, C-92
- VLAN, best route selection ... C-94

documentation

- feature matrix ... -xxvi
- latest versions ... -xxv
- printed in-box publications ... -xxv
- release notes ... -xxv

Domain Name Server

See DNS.

download

- software ... 15-17, A-25
- software using TFTP ... A-5
- switch-to-switch ... A-24
- TFTP ... A-6
- troubleshooting ... A-7
- Xmodem ... A-20

See also switch software.

duplex advertisements ... 14-56

duplex information, displaying ... 14-73

duplicate MAC address

See MAC address.

Dyn1

See LACP.

dynamic ARP protection

resource usage ... E-2

E

edge ports ... 13-4

Emergency Location Id Number ... 14-39, 14-67

erase

config file ... 6-39

event log

clearing entries ... C-36
compared to debug/Syslog operation ... C-41
console menu ... 3-7
debugging by severity level ... C-42, C-54
debugging by system module ... C-42, C-54
Event Log Message Reference Guide ... -xxv
format, date ... C-28
generated by system module ... C-28
how to read entries ... C-27
listing entries ... C-36
losing messages ... C-27
navigation ... C-35
not affected by debug configuration ... C-61
security levels ... 14-20
sending event log messages as traps ... 14-20
sending messages to Syslog server ... C-42
severity level ... C-27, C-59
system module ... C-60
time format ... C-28
UDLD warning messages ... 10-39
used for debugging ... C-42
used for troubleshooting ... C-27

excessive frames ... 13-40

F

fabric modules

disabling ... 15-12
enabling ... 15-12

facility

logging ... C-42

factory default configuration

restoring ... 6-9, C-83

failover, locator LED ... C-96

failover, management module, locator

LED ... C-96

failure, switch software download ... A-8

fans, show status ... B-8

fastboot command ... 6-24

fault detection policy ... 5-9, 5-24

fault-tolerance ... 12-5

fiber optics, monitoring links ... 10-33

filter, source-port

jumbo VLANs ... 13-38

firmware version ... B-7

flash memory ... 3-10, 6-3

flow control

constraints ... 10-5, 10-18
effect on rate-limiting ... 13-9, 13-16
global ... 10-18
global requirement ... 10-5
jumbo frames ... 13-38
per-port ... 10-5, 10-18
status ... B-15
terminal ... 7-3

flow sampling ... 14-4

friendly port names

See port names, friendly.

G

gateway

configuring ... 8-5
default gateway ... 8-3
IP address ... 8-4, 8-6
manual config priority ... 8-13
on primary VLAN ... 8-4
precedence of manual gateway over DHCP/
Bootp ... 8-13
routing fails ... C-12

giant frames ... 13-40

global config level ... 8-10

GMB

See guaranteed minimum bandwidth.

guaranteed minimum bandwidth

apportioning unallocated bandwidth ... 13-24
configuration ... 13-25
described ... 13-22
displaying current configuration ... 13-28
impacts of QoS queue configuration ... 13-24
operation ... 13-22
outbound queue priority ... 13-23

starving queues ... 13-24

H

Help

for CLI ... 1-7, 4-11
for menu interface ... 1-6, 3-9, 3-11
for web browser interface ... 1-7, 5-14
online, inoperable ... 5-14

hop, router ... 8-11

hotswapping mgmt module ... 15-15

HP

Auto-MDIX feature ... 10-21
web browser interface ... 2-5

I

ICMP

resources ... E-4

ICMP rate-limiting

all-traffic
 See rate-limiting.
caution ... 13-11
configuring ... 13-13
current rate-limit configuration ... 13-14
effect of flow control ... 13-16
effect on port trunks ... 13-15
effects of ... 13-12
event log messages ... 13-17
interface support ... 13-15
monitoring/mirroring ... 13-15
network application ... 13-12
no meshing ... 13-15
note on testing ... 13-17
operating notes ... 13-15
operation ... 13-10, 13-13
optimum packet size ... 13-16
resource usage ... E-2
spoofed ping ... 13-12
See also ICMP and rate-limiting.

IDM

resource usage ... E-2
resources ... E-4, E-7

IDS ... B-32

IEEE 802.1d ... C-19

IEEE P802.1AB/D9 ... 14-44

IGMP

host not receiving ... C-14

not working ... C-14

statistics ... B-24

inactivity timeout ... 7-4

inactivity-timer ... 7-9

Inbound Telnet Enabled parameter ... C-7

include-credentials, SNMP ... 9-23

informs

sending to trap receiver ... 14-20
SNMP ... 14-21

intelligent mirroring

See mirroring.

IP ... 8-7

address maximums ... G-1
CLI access ... 8-7
configuration ... 8-2
DHCP/Bootp ... 8-3
duplicate address ... C-8
duplicate address, DHCP network ... C-8
effect when address not used ... 8-11
features available with and without ... 8-12
gateway ... 8-3
gateway (IP) address ... 8-4
menu access ... 8-5
multiple addresses in VLAN ... 8-3, 8-9
single source addressing ... 8-25
source IP address ... 8-26
source-interface command ... 8-27
subnet ... 8-3, 8-9
subnet mask ... 8-2, 8-6
time server address ... 9-10, 9-29
Time-To-Live ... 8-7, 8-11
TTL ... 8-7, 8-11
using for web browser interface ... 5-5
web access ... 8-11

IP address

Configured IP address ... 8-27
Configured IP interface ... 8-27
displaying source IP information ... 8-29
displaying source-interface status ... 8-27
for SNMP management ... 14-3
loopback interface configuration ... 8-17
manually configure ... 8-6
multiple in a VLAN ... 8-9
outgoing interface ... 8-26
quick start ... 1-8, 8-4
removing or replacing ... 8-10
setup screen ... 8-4
show management command ... 8-7

- single source ... 8-25
- source IP address ... 8-26
- source IP with debug
 - debug
 - source IP address ... 8-31
- source IP with radius ... 8-31
- source IP with tacacs ... 8-31
- source-interface option ... 8-26

IP Preserve

- DHCP server ... 8-21
- overview ... 8-21
- rules, operating ... 8-21
- summary of effect ... 8-24

IP routing

- debug messages ... C-42

IPv6

- debug dhcpv6 messages ... C-43
- match criteria
 - classifier ... B-34, B-56, B-66

IPX

- broadcast traffic ... 10-5, 10-21
- network number ... B-11

J

jumbo frames

- configuration ... 13-32
- excessive inbound ... 13-38
- flow control ... 13-38
- GVRP operation ... 13-31
- management VLAN ... 13-37
- maximum size ... 13-30, 13-35
- meshing ... 13-31
- mirroring ... B-92
- MTU ... 13-30, B-92, B-95
- port adds and moves ... 13-31
- port speed ... 13-31
- security concerns ... 13-38
- standard MTU ... 13-31
- switch mesh domain ... 13-39
- through non-jumbo ports ... 13-39
- traffic sources ... 13-31
- troubleshooting ... 13-40
- VLAN tag ... 13-30, B-93
- voice VLAN ... 13-37

K

- kill command** ... 7-11

L

LACP

- 802.1X not allowed ... 12-23
- active ... 12-16
- blocked ports ... 12-24
- CLI access ... 12-12
- default port operation ... 12-22
- described ... 12-7, 12-19
- Dyn1 ... 12-8
- dynamic ... 12-20
- enabling dynamic trunk ... 12-16
- full-duplex required ... 12-5, 12-19
- IGMP ... 12-24
- mirroring static trunk ... B-36
- no half-duplex ... 12-26
- operation not allowed ... C-14
- overview of port mode settings ... 12-6
- passive ... 12-16
- removing port from active trunk ... 12-17
- restrictions ... 12-23
- standby link ... 12-20
- status, terms ... 12-22
- STP ... 12-24
- trunk limit ... 12-20
- VLANs ... 12-24
- with 802.1X ... 12-23
- with port security ... 12-23

- layer-3 scalability** ... G-1

- LED, locator** ... C-96

- licensing** ... H-1

- limit, broadcast** ... 10-20

- line rate** ... B-36

link failures

- detecting ... 10-33

- link speed, port trunk** ... 12-4

- link test** ... C-63

- link, serial** ... 7-3

- link-change traps** ... 14-17, 14-28

Link-Layer Discovery Protocol

- See* LLDP.

listening

- snmp-server ... 14-33

LLDP

- 802.1D-compliant switch ... 14-79

- 802.1X blocking ... 14-46
- 802.1X effect ... 14-79
- active port ... 14-39
- adjacent device ... 14-39
- advertisement ... 14-39
- advertisement content ... 14-54
- advertisement data ... 14-71
- advertisement, mandatory data ... 14-54
- advertisement, optional data ... 14-55
- advertisements, delay interval ... 14-50
- CDP neighbor data ... 14-80
- chassis ID ... 14-54
- chassis type ... 14-54
- clear statistics counters ... 14-76
- comparison with CDP data fields ... 14-81
- configuration options ... 14-42
- configuring optional data ... 14-55
- data options ... 14-43
- data read options ... 14-44
- data unit ... 14-40
- debug logging ... 14-44
- debug messages ... C-42, C-43
- default configuration ... 14-46
- DHCP/Bootp operation ... 14-45
- disable, per-port ... 14-53
- display neighbor data ... 14-74
- ELIN ... 14-39
- enable/disable, global ... 14-48
- features ... 14-38
- general operation ... 14-41
- global counters ... 14-76
- holdtime multiplier ... 14-50
- hub, packet-forwarding ... 14-41
- IEEE P802.1AB/D9 ... 14-44
- inconsistent value ... 14-51
- information options ... 14-43
- invalid frames ... 14-77
- IP address advertisement ... 14-45, 14-79
- IP address subelement ... 14-54
- IP address, DHCP/Bootp ... 14-54
- IP address, options ... 14-54
- IP address, version advertised ... 14-54
- LLDP-aware ... 14-39
- LLDPDU ... 14-40
- mandatory TLVs ... 14-79
- MIB ... 14-41, 14-44
- neighbor ... 14-40
- neighbor data remaining ... 14-79
- neighbor data, displaying ... 14-74
- neighbor statistics ... 14-76
- neighbor, maximum ... 14-78
- operating rules ... 14-45
- operation ... 14-41
- optional data, configuring ... 14-55
- outbound packet options ... 14-43
- packet boundaries ... 14-41
- packet dropped ... 14-41
- packet time-to-live ... 14-44
- packet-forwarding ... 14-41, 14-79
- packets not forwarded ... 14-40
- per-port counters ... 14-77
- port description ... 14-55
- port ID ... 14-54
- port speed ... 14-56
- port trunks ... 14-45
- port type ... 14-54
- refresh interval ... 14-49
- reinitialization delay ... 14-51
- remote management address ... 14-44
- remote manager address ... 14-54
- reset counters ... 14-76
- rxonly ... 14-53
- setmib, delay interval ... 14-50
- setmib, reinit delay ... 14-52
- show advertisement data ... 14-71
- show commands ... 14-46, 14-48
- show outbound advertisement ... 14-72
- SNMP notification ... 14-43
- SNMP traps ... 14-43
- spanning-tree blocking ... 14-46
- standards compatibility ... 14-44
- statistics ... 14-76
- statistics, displaying ... 14-76
- system capabilities ... 14-55
- system description ... 14-55
- system name ... 14-55
- terminology ... 14-39
- time-to-live ... 14-42, 14-50
- TLV ... 14-41
- transmission frequency ... 14-42
- transmission interval, change ... 14-49
- transmit and receive ... 14-42
- transmit/receive modes ... 14-42
- transmit/receive modes, per-port ... 14-53
- trap notice interval ... 14-53
- trap notification ... 14-52

- trap receiver, data change notice ... 14-52
- TTL ... 14-42, 14-44
- txonly ... 14-53
- VLAN, untagged ... 14-79
- walkmib ... 14-44
- with PoE ... 11-18

LLDP-MED

- displaying speed ... 14-73
- ELIN ... 14-67
- enable or disable ... 14-42
- endpoint support ... 14-58
- fast start control ... 14-62
- location data ... 14-66
- medTlvenable ... 14-64
- Neighbors MIB ... 14-74
- topology change notification ... 14-60
- Voice over IP ... 14-57

load balancing

- See* port trunk.

locator LED ... C-96

logging

- facility ... C-42

logging command ... C-50

- syntax ... C-42, C-54

logical port ... 12-9

loop, network ... 12-4

loopback interface

- benefits ... 8-16
- configuration ... 8-17
- default ... 8-16, 8-19
- displaying configuration ... 8-18
- in OSPF area ... 8-16
- multiple interfaces supported ... 8-16

lost password ... 5-11

M

MAC address ... 8-14, B-7, D-2

- displaying detected devices ... D-7
- duplicate ... C-19, C-25
- learned ... B-19
- per-slot or per-switch ... D-5
- port ... D-2, D-4
- same MAC, multiple VLANs ... D-6
- switch ... D-2
- traffic selection in mirroring ... B-34
- VLAN ... D-2, D-5
- walkmib ... D-5

MAC authentication

- SNMP notification ... 14-26

Maintenance Power Signature ... 11-3

management

- interfaces described ... 2-2
- server URL ... 5-13, 5-14
- server URL default ... 5-14

Management Information Base

- See* MIB.

management module failover, locator

- LED ... C-96**

management port ... J-2

management VLAN

- See* VLAN.

management VLAN, DNS ... C-94

manager access ... 4-5, 4-6, 14-13

manager password ... 5-9, 5-11

- SNMP notification ... 14-17, 14-26

manager privileges ... 4-5, 4-6

match criteria

- mirroring, classifier-based ... B-29

max frame size, jumbo ... 13-35

maximums ... G-1

MD5 authentication ... 14-9

MDI/MDI-X

- configuration, display ... 10-23
- operation ... 10-21
- port mode, display ... 10-23

media type, port trunk ... 12-4

memory

- flash ... 3-10, 6-3
- startup configuration ... 3-10

menu interface

- configuration changes, saving ... 3-10
- moving to or from the CLI ... 4-7

- See also* console.

mesh

- jumbo frames ... 13-39
- mirroring ... B-28

meshed ports, mirroring ... B-33

MIB

- HP proprietary ... 14-4
- listing ... 14-4
- standard ... 14-4

mini-GBICs, displaying info ... 15-23

mirroring

- 802.1Q tag ... B-95

- ACL criteria (deprecated) ... B-29, B-35, B-37, B-56, B-62
- ACLs converted to classifier-based policies in K.14.xx and later ... B-30, B-62
- ACLs replaced by classifier-based criteria ... B-29, B-67
- ARP request ... B-96
- booting pre-K.12.xx OS ... B-37
- caution
 - configure destination first ... B-33, B-44, B-53, B-71
 - endpoint removal ... B-51
 - exit port connection ... B-32, B-97
- classifier, first release ... B-29
- CLI option ... B-29, B-36
- command index ... B-43
- configuration options ... B-35
- configuration override ... B-39
- configuration, endpoint switch ... B-50
- configuration, Menu ... B-40
- configuration, session identity ... B-49
- configuration, source switch ... B-52
- configuration, traffic selection ... B-55
- destination
 - classifier-based configuration ... B-71
 - configuration ... B-78
 - local session ... B-33, B-43
 - remote session ... B-33, B-43
- display configuration ... B-76
- distributed traffic ... B-36
- dropped traffic ... B-36, B-95
- duplicate frames, IGMP ... B-95
- dynamic LACP trunk not supported ... B-36
- effect of STP state ... B-95
- encapsulation ... B-37
- encryption ... B-95
- endpoint configuration ... B-51, B-78
- endpoint switch ... B-36, B-47
- examples
 - classifier-based ... B-88
 - local session ... B-31, B-87
 - remote session ... B-31, B-90
- exit port
 - caution ... B-32
 - local session ... B-31, B-44
 - oversubscribe ... B-35, B-36
 - remote session ... B-31, B-48
 - requirements ... B-36
 - VLAN subnet ... B-31, B-34, B-46, B-48, B-49, B-90, B-97
- exit switch ... B-32
 - configuration ... B-78
- frame fragment ... B-37
- frame truncation, not supported ... B-37, B-92
- header ... B-29
- IGMP, duplicate frames ... B-95
- in show running configuration ... B-86
- intelligent mirroring ... B-28
- intermediate switches ... B-36
- intrusion detection system (IDS) ... B-28, B-32
- IPv4 encapsulation ... B-29, B-32, B-33, B-37, B-44, B-53, B-71, B-92, B-95
- IPv4 frames not mirrored ... B-96
- jumbo frames ... B-29, B-92
- K.12.xx, earlier software ... B-37
- legacy configuration ... B-37
- local session
 - configuration steps ... B-44
 - defined ... B-32
 - exit port ... B-31
 - quick reference ... B-45
- maximum sessions
 - destination ... B-28
 - source ... B-29
- maximum sources on a destination ... B-28
- Menu interface limit ... B-36, B-39
- Menu interface, local-only ... B-36
- Menu option ... B-29
- meshed ports ... B-33
- migration to K.12.xx ... B-37
- migration to release K.14.xx ... B-37
- mirror command ... B-53
- monitor, autoconfig session 1 ... B-57, B-58, B-61, B-96
- monitored interface
 - mesh ... B-35, B-41
 - of source traffic ... B-32, B-33
 - port ... B-35, B-41
 - trunk ... B-35, B-41
 - VLAN ... B-35, B-41
- MTU ... B-92, B-95
- no-tag-added ... B-59, B-95
- operating notes ... B-95
- overload on destination ... B-36
- oversized frames ... B-37
- port screen ... B-28

- rate ... B-36
- remote session
 - configuration steps ... B-46
 - defined ... B-32
 - disabling ... B-48
 - exit port ... B-31
 - first release supported ... B-28
 - quick reference ... B-47
 - supported switches ... B-33
- restrictions
 - classifier-based ... B-67, B-72, B-73
 - local sessions ... B-28
 - remote sessions ... B-28
 - source switch ... B-56
- session 1, legacy configuration ... B-37
- session limits ... B-34
- show commands ... B-76, B-78, B-79, B-86
- simultaneous source/destination ... B-34
- SNMP ... B-36, B-39
- SNMP for no-tag-added mirroring ... B-59
- source switch ... B-32
- static trunk ... B-33
- static VLAN ... B-33
- terminology ... B-30
- traffic overload ... B-36
- traffic selection
 - classifier-based criteria ... B-29, B-31, B-34, B-35, B-56, B-66, B-67, B-68, B-69, B-77
 - direction-based criteria ... B-29, B-31, B-34, B-57
 - MAC-based criteria ... B-29, B-34, B-35, B-56, B-63, B-77
 - overview ... B-29, B-34
- traffic, injected into mirrored stream ... B-96
- traffic, intercepted ... B-96
- troubleshooting ... B-97
- UDP destination address ... B-77, B-78, B-79
- UDP port ... B-77, B-78, B-79, B-90
- UDP source address ... B-78, B-79
- untagged mirror packets ... B-59
- VLAN
 - maximum frame size ... B-93
 - no-tag-added ... B-58, B-95
 - subnet, exit port ... B-31
 - tagged/untagged frames retained ... B-95
- Web interface ... B-36
- Web limits ... B-39

MLTS ... 14-40

module

- clearing the config ... 10-31
- CLI command ... 10-31
- configuring when not inserted ... 10-31
- pre-configuring ... 10-31
- remove configuration command ... 10-32

modules

- mini-GBIC information ... B-14
- show command ... B-13
- show details ... B-13

modules, show command ... 15-23

monitoring

- links between ports ... 10-33
- locator LED ... C-96
- status and counters screens ... B-5

monitoring, traffic
See mirroring.

MPS, defined ... 11-3

multicast ... 13-20

Multiline Telephone system ... 14-40

multinetting ... 8-3, 8-9
See also ACLs.

multiple configuration file
See configuration file, multiple.

multiple forwarding database ... B-11, B-22

multiple VLAN ... 14-3

N

NANP ... 14-40

navigation, event log ... C-35

network management functions ... 14-5, 14-13

network manager address ... 14-4, 14-5

network slow ... C-8

North American Numbering Plan ... 14-40

no-tag-added ... B-59, B-95

notifications

- authentication messages ... 14-17, 14-26
- configuring trap receivers ... 14-19
- enabling for network security ... 14-26
- link-change traps ... 14-17
- network security ... 14-26

O

online Help
See Help.

oobm

- address config ... J-10
- client commands ... J-14
- command ... J-6
- copy command output ... A-40
- copy config to remote host ... A-30
- copy crash-data ... A-42, A-43
- copy crash-log ... A-43, A-44
- copy event-log tftp ... A-41
- copy show-tech ... A-32
- copy tftp command-file ... A-36
- copy tftp config ... A-31
- copy tftp flash ... A-26
- default gateway config ... J-10
- enable/disable ... J-7, J-8
- server commands ... J-13
- show arp ... J-12
- show commands ... J-11
- show config ... J-12
- SNTP ... 9-13
- speed-duplex ... J-9
- telnet ... 7-7
- telnet-server ... 7-6
- tftp ... A-10
- tftp copy ... A-27, A-28
- tftp download ... A-8
- tftp traffic ... 6-38, 6-39
- timep ... 9-32
- transferring files ... 6-38
- web-management ... 7-8

operating system

See switch software.

operation not allowed, LACP ... C-14

operator access ... 4-4, 4-6, 14-13

operator password ... 5-11

setting via web browser ... 5-9

operator privileges ... 4-4, 4-6

Option 66, DHCP ... 6-41

OS

version ... A-25

See also switch software.

OSPF

- debug command ... C-51
- debug messages ... C-43
- using loopback interface as router ID ... 8-16

out-of-band access ... 2-3

P

packet

debug messages ... C-43

password ... 5-9, 5-11

- console ... 3-7
- creating ... 5-9
- delete ... 5-11
- disables usb autorun ... A-51
- if you lose the password ... 5-11
- lost ... 5-11
- manager ... 4-4, 5-9
- operator ... 4-4, 5-9
- setting ... 5-10
- SNMP notification ... 14-26
- SNMP notification for invalid login ... 14-17
- using to access browser and console ... 5-11
- web interface ... 5-9

pattern matching, show command output ... C-79

PD ... 14-40

ping ... C-87, C-89, C-92

See also DNS, resolver.

See also troubleshooting.

ping test ... C-63

PoE

- active ports, defined ... 11-3
- advertisements ... 14-66
- allocate-by ... 11-6
- allocation, usage ... 11-12
- benefit of LLDP-MED ... 14-58
- changing priority level ... 11-8
- changing the threshold ... 11-8, 11-16
- configuration options ... 11-5
- configuration planning ... 11-25
- configuring operation ... 11-8
- configuring port priority ... 11-9
- configuring redundancy ... 11-14
- detection status ... 11-21
- displaying power status ... 11-19
- enable or disable operation ... 11-5, 11-8
- enabling, disabling redundancy ... 11-14
- EPS, defined ... 11-3
- event log messages ... 11-28
- fault ... 11-13
- LLDP detection, enabling or disabling ... 11-18
- manually configuring power levels ... 11-13
- max module power ... 11-5, 11-7
- messages ... 11-28
- MPS

- absent cnt ... 11-24
- defined ... 11-3
- needed power for PoE+ ... 11-7
- other fault ... 11-23
- over current cnt ... 11-23
- oversubscribed ... 11-3
- overview of status ... 11-21
- PD support ... 11-6
- PD, defined ... 11-3
- poe-lldp-detect command ... 11-18
- port-number priority ... 11-7
- port-number priority, defined ... 11-4
- power denied cnt ... 11-23
- power, provisioning ... 11-5
- prioritizing power ... 11-7
- priority class ... 11-4, 11-7
 - defined ... 11-3
- priority critical ... 11-10
- priority high ... 11-10
- priority low ... 11-10
- priority policies ... 11-26
- priority, port ... 11-6, 11-7
- PSE, defined ... 11-3
- QoS classifiers ... 11-26
- RPS, defined ... 11-3
- security ... 11-26
- setting allocation ... 11-13
- short cnt ... 11-24
- slot-id-range option ... 11-16
- status ... 14-62
- status on specific ports ... 11-23
- supporting pre-standard devices ... 11-8
- terminology ... 11-3
- threshold, power ... 11-15
- usage ... 11-6
- using LLDP ... 11-18
- VLAN assignments ... 11-26
- policy enforcement engine**
 - described ... E-2
 - displaying resource usage ... E-2
- poll interval**
 - See* TimeP.
- port**
 - address table ... B-19
 - blocked by UDLD ... 10-34
 - broadcast limit ... 10-20
 - CLI access ... 10-8
 - configuration ... 10-1
 - configuring UDLD ... 10-34
 - context level ... 10-16
 - counters ... B-16
 - counters, reset ... B-16
 - default loopback interface ... 8-19
 - displaying loopback interface ... 8-18
 - duplex, view ... 10-8
 - enabling UDLD ... 10-35
 - fiber-optic ... 10-5
 - loopback interface configuration ... 8-16, 8-17
 - MAC address ... D-4, D-5
 - management ... J-2
 - menu access ... 10-6
 - mirroring, static LACP trunk ... B-36
 - monitoring ... B-28
 - speed, view ... 10-8
 - traffic patterns ... B-16
 - transceiver status ... 10-14
 - trunk
 - See* port trunk.
 - utilization ... 5-18, 10-13
 - CLI ... 10-13
 - web browser interface ... 5-18
 - web browser access ... 10-24
- port configuration** ... 12-1
- port mirroring**
 - See* mirroring.
- port names, friendly**
 - configuring ... 10-26
 - displaying ... 10-27
 - summary ... 10-25
- port security**
 - port trunk restriction ... 12-4
 - trunk restriction ... 12-9
- port trunk** ... 12-3
 - bandwidth capacity ... 12-3
 - caution ... 12-4, 12-10, 12-18
 - CLI access ... 12-12
 - default trunk type ... 12-11
 - enabling dynamic LACP ... 12-16
 - enabling UDLD ... 10-35
 - IGMP ... 12-9
 - limit ... 12-3
 - limit, combined ... 12-20
 - link requirements ... 12-4
 - logical port ... 12-9
 - media requirements ... 12-8
 - media type ... 12-4

- menu access to static trunk ... 12-10
- mirroring ... B-28
- monitor port restrictions ... 12-9
- nonconsecutive ports ... 12-3
- port security restriction ... 12-9
- removing port from static trunk ... 12-16
- requirements ... 12-8
- SA/DA ... 12-37
- spanning tree protocol ... 12-9
- static trunk ... 12-8
- static trunk, overview ... 12-6
- static/dynamic limit ... 12-20
- STP ... 12-9
- STP operation ... 12-8
- traffic distribution ... 12-8
- Trk1 ... 12-8
- trunk (non-protocol) option ... 12-7
- trunk option described ... 12-36
- types ... 12-7
- UDLD configuration ... 10-34
- VLAN ... 12-9
- VLAN operation ... 12-8
- web browser access ... 12-18
- See also* LACP.
- port trunk group**
 - interface access ... 12-1
- port, active** ... 14-39
- port-access authentication**
 - SNMP notification ... 14-26
- port-based access control**
 - event log ... C-15
 - LACP not allowed ... 12-23
 - troubleshooting ... C-15
- port-utilization and status displays** ... 10-13
- power levels, configuring** ... 11-13
- power supply**
 - show settings ... B-8
- power-over-ethernet**
 - See* PoE.
- Power-Sourcing Equipment** ... 11-4, 14-40
- Premium License**
 - installing ... H-1
 - overview, list of features ... -xxvi
- priority class**
 - defined ... 11-4
- priority of operation** ... 11-5
- privilege levels** ... 4-3
- ProCurve**

- Auto-MDIX feature ... 10-21
- support URL ... 5-14
- switch documentation ... -xxv
- ProCurve Manager**
 - reading USB autorun files ... A-49
 - required for USB autorun ... A-47
 - security concerns when deleting public community ... 14-5
 - SNMP and network management ... 14-3
 - starting web browser ... 5-5
 - updating switch software ... A-27
 - using Java-enabled browser ... 5-6
- ProCurve, HP, URL** ... 14-4
- prompt, =>** ... C-84
- PSAP** ... 14-40
- PSE** ... 14-40
- PSE, defined** ... 11-4
- Public Safety Answering Point** ... 14-40
- public SNMP community** ... 14-5, 14-13

Q

QoS

See Quality of Service.

Quality of Service

- queue configuration ... 13-24
- resource usage ... E-2
- resources ... E-4

quick configuration

... 3-8

quick start

... 1-8

R

RADIUS

- web browser access ... 5-9

RADIUS-assigned ACLs

- resources ... E-2

rate display for ports

... 10-13

rate-limiting

- bcast command ... 13-19
- broadcast traffic ... 13-19
- caution ... 13-4
- configuration ... 13-5, 13-13
- disabling multicast ... 13-21
- displaying configuration ... 13-6, 13-14
- edge ports ... 13-4
- effect of flow control ... 13-9, 13-16
- effect on port trunks ... 13-8, 13-15

- how measured ... 13-9
- ICMP
 - See ICMP rate-limiting.
- intended use ... 13-4
- mcast command ... 13-19
- multicast traffic ... 13-19
- note on testing ... 13-10, 13-17
- operating notes ... 13-8
- optimum packet size ... 13-10, 13-16
- per-port only ... 13-4
- purpose ... 13-4
- traffic filters ... 13-9
- reboot**
 - actions causing ... 6-4
 - faster boot time ... 6-24
 - from secondary flash ... 6-23
 - obtaining faster reboot time ... 6-20
 - scheduling remotely ... 6-25
 - via menu console ... 3-8
 - via menu interface ... 3-10, 3-12
 - See also* boot.
- redo, command description** ... 4-16
- redundancy** ... 11-14
 - boot command ... 15-29
 - boot-history ... 15-42
 - causes of switchover ... 15-13
 - disabling ... 15-6, 15-21
 - downloading software ... 15-17
 - enabling ... 15-6
 - event log messages ... 15-46
 - hotswapping module ... 15-15
 - how active module determined ... 15-44
 - locator LED ... C-96
 - log messages ... 15-41
 - reload command ... 15-32
 - resetting mgmt module ... 15-14
 - setting active module ... 15-9
 - setting default flash for boot ... 15-31
 - show flash ... 15-25
 - show log ... 15-26
 - show module ... 15-23
 - show redundancy ... 15-5, 15-24
 - show version ... 15-25
 - software version mismatch ... 15-16, 15-18
 - using web browser interface ... 15-36
- redundancy active-management** ... 15-9
- redundancy switchover** ... 15-8
- reload** ... 6-4, 15-32
- reload command** ... 6-20
- remote intelligent mirroring**
 - See* mirroring.
- remote mirroring**
 - resource usage ... E-2
- remote session, terminate** ... 7-11
- repeat, command description** ... 4-16
- Reset button** ... 6-4
 - restoring factory default configuration ... C-84
- reset operating system** ... 3-12
- reset port counters** ... B-16
- resetting the switch**
 - factory default reset ... C-83
- resource monitor**
 - event log ... E-7
- resource usage**
 - displaying ... E-4
 - insufficient resources ... E-7
- restricted write access** ... 14-13
- RFCs**
 - RFC 1493 ... 14-4
 - RFC 1515 ... 14-4
 - RFC 2737 ... 14-44, 14-45
 - RFC 2863 ... 14-44, 14-45
 - RFC 2922 ... 14-44
 - RFC 3176 ... 14-34
 - See also* MIB.
- RIP**
 - broadcast traffic ... 10-5, 10-21
 - debug command ... C-51
 - debug messages ... C-43
- RMON** ... 14-4
- RMON groups supported** ... 14-34
- router**
 - gateway ... 8-6
 - maximum routes ... G-1
 - OSPF area maximum ... G-2
 - OSPF interface maximum ... G-2
 - RIP interface maximum ... G-2
 - supported routes ... G-1
- router, hop** ... 8-11
- routing**
 - gateway fails ... C-12
 - OSPF debug messages ... C-51
 - RIP debug messages ... C-51
 - traceroute ... C-67
- RS-232** ... 2-3
- running-config**

viewing ... 6-6
See also configuration.

S

savepower

command ... I-3
led option ... I-4
port-low-pwr ... I-6
show led ... I-7
show module ... I-6
show port-low-pwr ... I-7

scalability ... G-1

scheduled reboot ... 6-25

SCP/SFTP

enabling ... A-13
session limit ... A-17, A-19
transfer of config files ... 6-40
troubleshooting ... A-18

secure copy

See SCP/SFTP.

secure FTP

See SCP/SFTP.

secure management VLAN

See VLAN.

secure management VLAN, DNS ... C-94

security

Clear button ... 5-12
enabling network security notifications ... 14-26
privilege levels in CLI ... 4-3
USB autorun ... A-48
username and password ... 5-9
web browser access, RADIUS ... 5-9

Self Test LED

behavior during factory default reset ... C-84

serial number ... B-7

setmib, delay interval ... 14-50

setmib, reinit delay ... 14-52

setup screen ... 1-8

severity level

event log ... C-27
selecting Event Log messages for
debugging ... C-59

sFlow ... 14-4

agent ... 14-34
CLI-owned versus SNMP-owned
configurations ... 14-35
configuring via the CLI ... 14-35

destination ... 14-34
sampling-polling information ... 14-37
show commands ... 14-35

SHA authentication ... 14-9

show

custom option ... 10-10
displaying specific output ... C-79
exclude option
show

begin option ... C-79

include option ... C-79
interfaces brief ... 10-8
interfaces config ... 10-9
pattern matching with ... C-79
tech, custom ... A-32
telnet ... 7-7

show cpu ... B-9

show debug ... C-46

show flash ... 15-25

show interfaces

dynamic display ... 10-9

show interfaces display ... C-78

show log ... 15-26

show management ... 9-10, 9-29

show module ... 15-23

show modules command ... 15-23

show modules, command ... B-13

show policy resources command ... B-84

show redundancy ... 15-24

show system

chassislocate ... B-8

show tech ... C-72

custom ... A-32

show version ... 15-25

show-tech ... A-31

slow network ... C-8

SNMP ... 14-3

ARP protection events ... 14-17
authentication notification ... 14-17, 14-26
CLI commands ... 14-13
communities ... 14-4, 14-5, 14-13, 14-14
configuring with the CLI ... 14-15
configuring with the menu ... 14-13
mapping ... 14-11
configure ... 14-4, 14-5
configuring security groups ... 14-23
configuring SNMPv3 notification ... 14-23
configuring SNMPv3 users ... 14-23

- configuring trap receivers ... 14-19
- configuring trap receivers ... 14-19
- DHCP snooping events ... 14-17
- different versions ... 14-17
- enabling informs ... 14-21
- enabling network security traps ... 14-27
- enabling SNMPv3 ... 14-23
- fixed traps ... 14-19
- invalid password in login ... 14-17
- IP ... 14-3
- link-change traps ... 14-17, 14-28
- manager password change ... 14-17
- mirroring ... B-36
- network security notification ... 14-26
- no-tag-added mirroring ... B-59
- notification, LLDP
 - SNMP notification ... 14-43
- public community ... 14-5, 14-13
- supported notifications ... 14-17
- system thresholds ... 14-19
- traps ... 10-34, 14-4, 14-17
- walkmib ... D-5, D-6
- well-known traps ... 14-19
- SNMP trap, LLDP** ... 14-52
- snmp-server**
 - listening mode ... 14-33
- SNMPv3**
 - "public" community access caution ... 14-6
 - access ... 14-5
 - assigning users to groups ... 14-7
 - authentication, configuring ... 14-9
 - communities ... 14-11
 - enable command ... 14-7
 - enabling ... 14-6
 - encryption, configuring ... 14-9
 - group access levels ... 14-11
 - groups ... 14-10
 - network management problems with snmpv3
 - only ... 14-6
 - restricted-access option ... 14-6
 - set up ... 14-5
 - users ... 14-5
- SNTP**
 - authentication command ... 9-20
 - authentication mode ... 9-18
 - broadcast mode ... 9-3, 9-12, 9-21
 - broadcast mode, requirement ... 9-4
 - client authentication ... 9-16
 - configuration ... 9-5
 - disabling ... 9-13
 - display config information ... 9-21
 - display statistics ... 9-22
 - enabling and disabling ... 9-11
 - event log messages ... 9-36
 - include-credentials ... 9-23, 9-24, 9-25
 - key-id ... 9-17, 9-18, 9-19, 9-20
 - key-value ... 9-17, 9-18
 - manual config priority ... 8-13
 - menu interface operation ... 9-36
 - oobm ... 9-13
 - operating modes ... 9-3
 - poll interval
 - See* TimeP.
 - priority ... 9-15, 9-20
 - selecting ... 9-4
 - server priority ... 9-15
 - show authentication ... 9-22
 - show management ... 9-10
 - trusted key ... 9-19
 - unicast mode ... 9-4, 9-12, 9-21
 - unicast time polling ... 9-35
 - unicast, deleting addresses ... 9-36
 - unicast, replacing servers ... 9-36
 - viewing ... 9-5, 9-9
- software**
 - See* switch software.
- software image**
 - See* switch software.
- software licensing** ... H-1
- software version** ... B-7
- sorting alert log entries** ... 5-21
- source port filters**
 - jumbo VLANs ... 13-38
- spanning tree**
 - fast-uplink, troubleshooting ... C-20
 - mirroring blocked traffic ... B-95
 - problems related to ... C-19
 - show tech, copy output ... C-73
 - using with port trunking ... 12-9
- SSH**
 - enabling or disabling ... A-15
 - file transfer ... A-11
 - TACACS exclusion ... A-16
 - troubleshooting ... A-18, C-20
- standard MIB** ... 14-4
- starting a console session** ... 3-4

- startup-config**
 - viewing ... 6-6
 - See also* configuration.
- statistics** ... 3-7
 - clearing ... B-18
 - SNTP ... 9-22
- statistics, clear counters** ... 6-11
- status and counters**
 - access from console ... 3-7
- status overview screen** ... 5-7
- subnet** ... 8-9
 - VLAN, mirroring exit port ... B-34, B-46, B-48, B-49, B-90, B-97
- subnet mask** ... 8-5, 8-6
 - See also* IP masks.
- support**
 - changing default URL ... 5-14
 - URL ... 5-13
 - URL Window ... 5-13
- switch console**
 - See* console.
- switch setup menu** ... 3-8
- switch software**
 - copy from a USB device ... A-22
 - download using TFTP ... A-5
 - download, failure indication ... A-8
 - download, switch-to-switch ... A-25
 - download, troubleshooting ... A-7
 - download, using TFTP ... A-5
 - installing a license ... H-1
 - software image ... A-4
 - version ... A-7, A-21
- switchover** ... 15-13
- Syslog**
 - "debug" severity level as default ... C-59, C-61
 - adding priority description ... C-58
 - compared to event log ... C-41
 - config friendly descriptions ... C-57
 - configuring for debugging ... C-44
 - configuring server address ... C-42
 - configuring server IP address ... C-50
 - configuring Syslog servers and debug destinations ... C-42
 - control-desc ... C-58
 - displaying Syslog configuration ... C-46
 - event log messages sent by default ... C-56
 - logging command ... C-50, C-52
 - operating notes ... C-60
 - overview ... C-41
 - priority-descr ... C-58
 - See also* debug command.
 - sending event log messages ... C-41
 - server configuration ... C-55
 - severity, "debug" ... C-52
 - specifying severity level events for debugging ... C-59
 - specifying system module events for debugging ... C-60
 - user facility as default ... C-57, C-61
 - using event log for debugging ... C-42, C-54
- system**
 - chassislocate ... B-8
- system configuration screen** ... 7-12
- system information** ... B-8
 - fans ... B-8
 - power-supply ... B-8
 - temperature ... B-8
- system module**
 - selecting event log messages for debugging ... C-60
- System Name parameter** ... 7-13
- System Support Modules, information about** ... 15-23

T

- TACACS**
 - SSH exclusion ... A-16
- task monitor** ... B-9
- taskusage -d** ... B-9
- taskUsageShow** ... B-9
- Telnet**
 - connecting to switch ... 3-4
 - enable/disable ... 7-4
 - outbound ... 7-6
 - terminate session, kill command ... 7-11
 - troubleshooting access ... C-7
- telnet**
 - domain name address ... 7-6
 - hostname ... 7-6
 - ipv6 address ... 7-6
 - oobm ... 7-7
 - show command ... 7-7
 - switch-num ... 7-6
- temperature, show settings** ... B-8
- terminal access, lose connectivity** ... 7-9

- terminal type** ... 7-3
- terminate remote session** ... 7-11
- TFTP**
 - auto-TFTP ... A-11
 - auto-TFTP feature ... A-11
 - auto-TFTP, disable ... A-11, A-14
 - copy command output ... A-40
 - copy crash data ... A-41
 - copy crash log ... A-43
 - copy event log output ... A-41
 - copying a configuration file ... A-31
 - copying software image ... A-28
 - disable ... A-14
 - disabled ... A-11
 - download software using CLI ... A-8
 - downloading software using console ... A-6
 - enable client or server ... A-10
 - enabling client functionality ... A-10
 - enabling server functionality ... A-10
 - switch-to-switch transfer ... A-24
 - troubleshooting download failures ... A-7
 - uploading an ACL command file ... A-36
 - using to download switch software ... A-5
- tftp**
 - listen mode ... A-10
 - oobm ... A-10, A-26, A-27, A-28
- threshold setting** ... 14-5, 14-13
- thresholds, SNMP** ... 14-19
- throttling, broadcast/multicast traffic** ... 13-19
- time format, events** ... C-28
- time protocol**
 - selecting ... 9-4
- time server** ... 8-2
- time zone** ... 7-13, 7-17
- time, configure** ... 7-18
- TimeP** ... 8-3, 8-5
 - assignment methods ... 9-3
 - disabling ... 9-33
 - enabling and disabling ... 9-30
 - manual config priority ... 8-13
 - poll interval ... 9-33
 - selecting ... 9-4
 - server address listing ... 9-10, 9-29
 - show management ... 9-29
 - viewing and configuring, menu ... 9-27
 - viewing, CLI ... 9-29
- timesync, disabling** ... 9-33
- Time-To-Live** ... 8-3, 8-5, 8-6, 8-11
 - See also* TTL.
- time-to-live, LLDP** ... 14-42
- Time-To-Live, on primary VLAN** ... 8-4
- TLV** ... 14-41
- TLVs, mandatory** ... 14-79
- traceroute** ... C-87, C-89, C-92
 - asterisk ... C-69
 - blocked route ... C-70
 - fails ... C-68
- traffic**
 - broadcast rate-limiting ... 13-19
 - multicast rate-limiting ... 13-19
- traffic mirroring**
 - See* mirroring.
- traffic monitoring** ... 14-5, 14-13
 - See also* sFlow and RMON.
 - See* mirroring.
- traffic, port** ... B-16
- transceiver**
 - error messages ... 10-15
 - view status ... 10-14
- transceiver, fiber-optic** ... 10-5
- transceivers**
 - configuring when not inserted ... 10-31
 - not inserted ... 10-31
- trap** ... 5-25
 - CLI access ... 14-19
 - configuring trap receivers ... 14-19
 - security levels ... 14-20
- trap notification** ... 14-52
- trap receiver** ... 14-4, 14-5
 - configuring ... 14-19
 - sending event log messages ... 14-20
 - sending SNMPv2 informs ... 14-20
 - SNMP ... 14-19
 - up to ten supported ... 14-19
- traps**
 - arp-protect ... 14-27
 - authentication trap ... 14-27
 - auth-server-fail ... 14-27
 - dhcp-snooping ... 14-27
 - dynamic-ip-lockdown ... 14-27
 - enabling network security notifications ... 14-26
 - fixed ... 14-19
 - link-change ... 14-27, 14-28
 - login-failure-mgr ... 14-27
 - password-change-mgr ... 14-27
 - port-security ... 14-27

See also notification.

- snmp-authentication ... 14-27
- threshold ... 14-19

troubleshooting

- ACL ... C-9
- approaches ... C-5
- browsing the configuration file ... C-71
- configuring debug destinations ... C-42
- console access problems ... C-6
- diagnosing unusual network activity ... C-8
- diagnostics tools ... C-62
- displaying switch operation ... C-72, C-75
- DNS
 - See* DNS.
- fast-uplink ... C-19
- ping and link tests ... C-63
- resource usage ... E-2
- restoring factory default configuration ... C-83
- spanning tree ... C-19
- SSH ... C-20
- SSH, SFTP, and SCP Operations ... A-18
- switch software download ... A-7
- switch won't reboot, shows => prompt ... C-84
- traceroute ... C-87, C-89
- unusual network activity ... C-8
- using CLI session ... C-42
- using debug and Syslog messaging
- using the event log ... C-27
- viewing switch operation ... C-71
- web browser access problems ... C-6

trunk

- distributed ... 12-27
- number supported ... 12-5
- See* port trunk.

TTL ... 8-3, 8-5, 8-6, 8-7

- IP ... 8-11
- LLDP ... 14-42
- manual config priority ... 8-13
- on primary VLAN ... 8-4
- See also* Time-To-Live.

Type-Length-Value ... 14-41

types of alert log entries ... 5-22

U

UDLD

- changing the keepalive interval ... 10-36
- changing the keepalive retries ... 10-36

- configuration ... 10-34
- configuring for tagged ports ... 10-36
- enabling on a port ... 10-35
- event log messages ... 10-39
- operation ... 10-34
- overview ... 10-33
- supported switches ... 10-34
- viewing configuration ... 10-37
- viewing statistics ... 10-38
- warning messages ... 10-39

unauthorized access ... 14-27

undersize frames ... 13-40

unicast mode

- SNTP ... 9-21

Uni-directional Link Detection

- See* UDLD.

Universal Resource Locator

- See* URL.

Unix, Bootp ... 8-14

unrestricted write access ... 14-13

unusual network activity ... C-8

up time ... B-7

URL

- browser interface online help location ... 5-14
- management ... 5-14
- management server ... 5-13, 5-14
- ProCurve ... 5-14, 14-4
- support ... 5-13, 5-14

USB

- autorun ... A-47–A-52
 - AutoRun file ... A-47
 - command file ... A-47
 - configuring passwords ... A-51
 - creating a command file ... A-47
 - enabling or disabling ... A-51
 - LED indications ... A-49
 - report outputs ... A-49
 - required software versions ... A-47
 - secure-mode ... A-51
 - security ... A-48
 - troubleshooting ... A-49
 - viewing config information ... A-52
- auxiliary port ... A-22, A-47
- auxiliary port LEDs ... A-49
- copy command output ... A-40
- copy configuration file to/from a USB
 - device ... A-35
- copy crash data ... A-41

- copy crash log ... A-43
- copy event log output ... A-41
- copy software image to a USB device ... A-29
- devices with secure partitions not supported ... A-22
- flash drives must be formatted ... A-22
- supported capabilities ... A-22
- uploading an ACL command file ... A-38
- using to copy switch software ... A-22
- viewing flash drive contents ... A-22, A-23

usb

- enable port ... 10-17

usb configuration ... 10-17

usb-port ... 10-17

user name

- using for browser or console access ... 5-9, 5-11

users, SNMPv3

- See* SNMPv3.

utilization, port ... 5-18, 10-13

V

version, OS ... A-25

version, switch software ... A-7, A-21

view

- duplex ... 10-8
- port speed ... 10-8
- transceiver status ... 10-14

virtual interface

- See* loopback interface

virus-throttling

- See* connection-rate filtering.

VLAN

- address ... 14-3
- Bootp ... 8-14
- configuring Bootp ... 8-14
- configuring UDLD for tagged ports ... 10-36
- device not seen ... C-24
- event log entries ... C-28
- ID ... 4-15
- IP address maximum ... G-1
- IP addressing with multiple ... 8-4
- jumbo max frame size ... 13-35
- link blocked ... C-19
- MAC address ... D-2, D-5
- management and jumbo frames ... 13-37
- management VLAN, resource usage ... E-2
- management VLAN, SNMP block ... 14-3

- maximums ... G-1
- mirroring ... B-4, B-28
- multinet ... 8-3
- multinetting ... 8-3, 8-9
- multiple ... 14-3
- multiple IP addresses ... 8-3, 8-9
- port configuration ... C-24
- primary ... 8-3
- reboot required ... 3-8
- same MAC, multiple VLANs ... D-6
- secure management VLAN, with DNS ... C-94
- subnet ... 8-3, 8-9
- support enable/disable ... 3-8
- switch software download ... A-5
- tagging broadcast, multicast, and unicast traffic ... C-24

VLAN ID

- See* VLAN.

VoIP

- LLDP-MED support ... 14-57

VRRP

- debug messages ... C-43

VT-100 terminal ... 7-3

W

walkmib ... 14-44, D-5, D-6

warranty ... -ii

web agent

- advantages ... 2-5
- disabling access ... 5-3
- enable/disable ... 7-4
- enabled parameter ... 5-3

Web authentication

- SNMP notification ... 14-26

web browser interface

- access configuration ... 7-3
- access parameters ... 5-9
- access security ... 7-3
- alert log ... 5-21
- alert log details ... 5-22
- bandwidth adjustment ... 5-19
- bar graph adjustment ... 5-19
- disable access ... 5-3
- enabling ... 5-5
- error packets ... 5-18
- fault detection policy ... 5-9, 5-24
- fault detection window ... 5-24

- features ... 2-5
- first-time install ... 5-8
- first-time tasks ... 5-8
- Java applets, enabling ... 5-5
- main screen ... 5-17
- online help ... 5-14
- online help location specifying ... 5-14
- online help, inoperable ... 5-14
- overview ... 5-17
- Overview window ... 5-17
- password lost ... 5-11
- password, setting ... 5-10
- port status ... 5-20
- port utilization ... 5-18
- port utilization and status displays ... 5-18
- screen elements ... 5-17
- security ... 5-3, 5-9
- standalone ... 5-5
- status indicators ... 5-23
- status overview screen ... 5-7
- system requirements ... 5-5
- troubleshooting access problems ... C-6
- URL default ... 5-14
- URL, management server ... 5-15
- URL, support ... 5-15
- using for redundant management ... 15-36
- web site, HP** ... 14-4
- web-management**
 - listen, oobm ... 7-8
 - oobm ... 7-8
- windshell, debug destination** ... C-42
- wireless services**
 - debug messages ... C-43
- world wide web site, HP**
 - See* ProCurve.
- write access** ... 14-13
- write memory**
 - effect on menu interface ... 3-13
 - redundant management ... 6-7
- copying a software image ... A-28
- download to primary or secondary flash ... A-21
- uploading an ACL command file ... A-38
- using to download switch software ... A-20

X

Xmodem

- copy command output ... A-40
- copy crash data ... A-41
- copy crash log ... A-43
- copy event log output ... A-41
- copying a configuration file ... A-33

Technology for better business outcomes

To learn more, visit www.hp.com/go/procurve/

© Copyright 2009 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. 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 will not be liable for technical or editorial errors or omissions contained herein.



5992-3059, September 2009