

Release Notes

Version 08.0.01k Operating System

for the ProCurve 9304M, 9308M, and 9315M Routing Switches
with Redundant Management (M2, M4, EP, and T-Flow), May 2007



Software release 08.0.01k is the follow-on to version 08.0.01c, which in turn superseded version 07.8.01d. (For information about software branches and minimum release requirements for management module support, refer to "Software Branches" on page 3.)

These release notes provide information on the following items:

- New hardware and software enhancements introduced since software version 07.8.01
- Known issues in this software release
- Procedure for upgrading the software code on ProCurve 9304M, 9308M, and 9315M Routing Switches with M2, M4, or EP Redundant Management modules. See "Upgrading Software on an M2, M4, or EP Management Module to Release 08.0.01k" on page 7.
- General procedures, usage information, and helpful notes for operating and managing ProCurve routing switches
- Software fixes since software version 07.0.01

Descriptions of the enhancements in release 07.8.00 are included in the manuals for the 07.8.00 release. If you purchased a Redundant Management module with software version 07.8.00 or greater installed, the CD shipped with the module includes the 07.8.00 manuals.

If you need to access ProCurve product documentation, refer to "Downloading the Latest Software and Documentation" on page 3 for information on how to download PDF versions of the latest manuals.

NOTES:

Software Update Notice: Check the ProCurve Website frequently for free software updates for various ProCurve switch products. (Refer to "Downloading the Latest Software and Documentation" on page 3.)

Mini-GBIC ports: Hewlett-Packard offers and supports only mini-GBICs having a ProCurve label (with product number J4858B, J4859B, or J4860B) for use with the following modules:

- J4856A ProCurve 9300 Mini-GBIC Module
- J4857A ProCurve 9300 Mini-GBIC Redundant Management Module
- J4885A ProCurve 9300 EP 8-port Mini-GBIC Redundant Management Module
- J4894A ProCurve 9300 EP 16-port Mini-GBIC Module
- J8177B ProCurve Gigabit-copper mini-GBIC is supported for use in only the EP modules (J4885A and J4894A)

Use of other brands of mini-GBICs is not supported.

Flash Images: The flash image files for this software release differ depending on the type of management module you use. Refer to "Boot Code Requirements for ProCurve Software" on page 5.

SNMP: Starting with software release 05.2.16, the software does not have a default read-write SNMP community. If you use the default community name "private" as the password for web management access or for read-write access through a network management application, you need to use the CLI to add the read-write community string first.

Devices Without Redundant Management: For information about how to upgrade software on the ProCurve 9304M and ProCurve 9308M routing switches without redundant management, refer to the latest 6.6.x release notes. (See "Downloading the Latest Software and Documentation" on page 3.)

© Copyright 2001, 2007 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice.

Publication Number

5991-4755
May 2007

Applicable ProCurve 9300 (Current) Products

9304M Routing Switch	(J4139A)
9308M Routing Switch	(J4138A)
9315M Routing Switch	(J4874A)
EP 10/100-TX RJ-45 Module	(J4881B)
EP 10/100-TX Telco (RJ-21) Module	(J4889B)
EP Mini-GBIC RM Module	(J4885A)
EP Mini-GBIC Module.	(J4894A)
EP 100/1000-T Module.	(J4895A)
2-Port 10 Gigabit Ethernet Module	(J8174A)
EP 100Base-FX Module	(J8178A)
Gigabit-SX-LC Mini-GBIC	(J4858B)
Gigabit-LX-LC Mini-GBIC	(J4859B)
Gigabit-LH-LC Mini-GBIC	(J4860B)
1000Base-T Mini-GBIC.	(J8177B)
10 Gigabit Ethernet LR Optic	(J8173A)
10 Gigabit Ethernet SR Optic	(J8175A)
10 Gigabit Ethernet ER Optic	(J8176A)
9304M/9308M Redundant Power Supply.	(J4147A)
9315M Redundant Power Supply.	(J4875A)

Trademark Credits

Microsoft®, Windows®, and Windows NT® are trademarks of Microsoft Corporation. Adobe® and Acrobat® are trademarks of Adobe Systems Incorporated.

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.

Hewlett-Packard Company
8000 Foothills Blvd.
Roseville, CA 95747-5551
USA
<http://www.procurve.com>

Contents

Terminology	1
Product Documentation for Software Release 08.0.01k	1
Downloading the Latest Software and Documentation	3
Software Branches	3
Software Requirements for Management Modules	4
Software Supported on ProCurve Routing Switches	5
Boot Code Requirements for ProCurve Software	5
Upgrading Software on an M2, M4, or EP Management Module to Release 08.0.01k	7
Restrictions	7
Upgrading to Software Release 08.0.01k: Overview	8
A. Upgrading Boot Code on a Management Module to Version 07.06.05	8
B. Upgrading Software on a Management Module From a Release Earlier than 07.6.01b	8
C. Upgrading Software on a Management Module From Release 07.6.01b or Greater	9
Using SNMP to Upgrade Software on a Management Module	10
Upgrading the FPGA on a 10 Gigabit Ethernet Module	11
Using Different Combinations of Management Modules	12
Non-Redundant Management on ProCurve 9304M and ProCurve 9308M Routing Switches	13
Maximum Size of Startup-Config and Running-Config Files	13
Removing a Module from an Active Device	14
Configuring the ProCurve 9315M	15
Minimum Software Release Supported	15
Inserting or Removing an EP Module on a ProCurve 9315M	15
Slot Locations for Redundant Management Modules	15
MAC Addresses	15
Server Trunk Groups	15
VLANs	16
Summary of Enhancements in 08.0.00	16
Layer 2 Enhancements in 08.0.00	16
Layer 3 Enhancements in 08.0.00	16
System-Level Enhancements in 08.0.00	17
Multicast Enhancements in 08.0.00	18
Security Enhancements in 08.0.00	19
Enhancements and Configuration Notes in 08.0.00	19
Clearing OSPF Information from the ProCurve Device	19
OSPF Redistribution Filter Rebinding	21
Configuring an OSPF Non-Broadcast Interface	21
Configuring a Unicast Route With Multiple Outgoing Ports	22
Advertising an IBGP Next Hop as a null0 Route as a Defense Against DDoS Attacks	23
IP Load Sharing for RIPv2 Routes	25
VRRPE Slow Start Timer	25
Dynamic Configuration of a Voice over IP (VoIP) phone	26
EP Layer 2 ACLs	27
Layer 2 ACL-Based Rate Limiting	30
802.1s Spanning Tree Support	32
Creating an Alias for a CLI Command	41

Directing Debugging Output to Multiple Destinations	42
sFlow Enhancements	43
Jumbo Packet Counter in show Command Output	45
Displaying VLANs in Numerical Order	46
Specifying a Host Name in an ACL Statement	46
Configuring a Domain Name List and Using Domain Look Up	47
IGMP V3 Snooping	54
Increased Size of the MSDP Source Active Cache	63
Specifying a Designated Router Election Priority for PIM V2	63
Support for Standard Multicast MIBs	64
Using Multi-Device Port Authentication and 802.1X Security on the Same Port	69
CPU ACL	75
Where to Get More Information	77
Software Fixes	78
Known Issues in 08.0.01k	100
ProCurve 9300M Series Modules	100

Terminology

The following table defines basic product terms used in ProCurve routing switch documentation.

Term	Definition
chassis device or chassis	A routing switch that accepts optional modules or power supplies. The ProCurve 9315M, ProCurve 9304M, and ProCurve 9308M routing switches are chassis devices.
EP (Enhanced Performance) and Standard	Routing switches can be EP or Standard devices, depending on whether the management module is an EP or Standard (M2 or M4) module. For a listing of ProCurve routing switches and their product numbers, see Table 4 on page 5.
routing switch or router	A Layer 3 device that switches and routes network traffic. The term <i>router</i> is sometimes used in this document in descriptions of a routing switch's Layer 3 routing protocol features.
switch	A Layer 2 device that switches network traffic.
ProCurve 9300#	An example Command Line Interface (CLI) prompt. Actual prompts show the product number for the routing switch, such as ProCurve 9300# .

Product Documentation for Software Release 08.0.01k

Software release 08.0.01k includes all of the features in release 07.8.00a, plus several new features.

- For documentation on the features in 08.0.01k that were available since 07.8.00a, refer to the product documentation set identified for "Software Version 07.8.00a".

Table 1 describes the main topics covered in the ProCurve Routing Switch documentation set. If you do not already have a PDF version of the 07.8.00a documentation set, refer to "Downloading the Latest Software and Documentation" on page 3.

Table 1: Where To Get More Information

Title	Contents
Installation and Basic Configuration Guide	<ul style="list-style-type: none"> • Installation • Basic Features <ul style="list-style-type: none"> • System (SNMP, SNTP, Syslog, broadcast and multicast throttling) • Port configuration (speed, mode) • Layer 2 (MAC table parameters, MAC filters, broadcast and multicast filters, port locks) • Parameter table resizing • Port monitoring • Link Aggregation • Spanning Tree Protocol • Virtual LANs • Layer 2 Multicast • Base Layer 3 • Upgrading Software (Important: See also “Upgrading Software on an M2, M4, or EP Management Module to Release 08.0.01k” on page 7 in this document.) • Hardware Specifications and RFCs
<i>Security Guide</i>	<ul style="list-style-type: none"> • Security (passwords, user accounts, AAA, RADIUS, and TACACS/TACACS+) • Secure Shell (SSH) • Denial of Service Protection
<i>Advanced Configuration and Management Guide</i>	<ul style="list-style-type: none"> • QoS • ACLs • EP rate limiting • Standard rate limiting • IP • RIP • IP Multicast • OSPF • BGP4 • Network Address Translation • VRRP and VRRPE • IPX • AppleTalk
Command Line Interface Reference	Syntax information for all CLI commands.
Diagnostics Guide	<ul style="list-style-type: none"> • Diagnostic commands • Backplane debugging commands • Changing CAM partitions

Downloading the Latest Software and Documentation

You can download the current software version and the latest routing switch product documentation from the ProCurve website as described below.

To Download a Software Version:

1. Go to the ProCurve website at <http://www.procurve.com>.
2. Click on **software** (in the sidebar).
3. Under “latest software”, click on **switches**.

NOTE: If you are downloading software for the ProCurve 9304M or ProCurve 9308M, select the option that matches the type of management module(s) you are using in the routing switch—with redundant management or without redundant management.

To Download Product Documentation:

For the latest version of product documentation for ProCurve routing switches:

1. Go to the ProCurve website at <http://www.procurve.com>.
2. Click on **technical support**, then **manuals**.
3. Click on the name of the product for which you want manuals.
4. On the page listing the manuals, find the latest manuals under the heading “**For software version 7.8.00a**”.

You will need the Adobe® Acrobat® Reader (version 4.0 or greater) to view or print the manuals.

Software Branches

Starting with software releases 06.6.28 and 07.1.10, ProCurve offers the software branches described in Table 2:

Table 2: Software Branches

Software Release	Includes:	Operates on:
06.6.28 and later 06.x releases	Bug fixes	ProCurve 9304M and ProCurve 9308M routing switches without redundant management (that is, with M1 modules) ProCurve 6308M-SX routing switch ProCurve 6208M-SX switch
07.1.10 and later 07.1.x releases	Bug fixes, new features, and enhancements to existing features	ProCurve 9304M and ProCurve 9308M routing switches with redundant management (M2 modules)
07.5.04 release	Bug fixes, new features, and enhancements to existing features	ProCurve 9304M, ProCurve 9308M, and ProCurve 9315M routing switches with redundant management (M2 and M4 modules)
07.6.00 and 07.6.01b releases	Bug fixes, new features, and enhancements to existing features	ProCurve 9304M, ProCurve 9308M, and ProCurve 9315M routing switches with redundant management (M2, M4, EP, and T-Flow) and 1-port 10GB modules
07.6.04 release 07.7.01 release	Bug fixes, new features, and enhancements to existing features, including support for the 2-port 10GB module	ProCurve 9304M, ProCurve 9308M, and ProCurve 9315M routing switches with redundant management (M2, M4, EP, and T-Flow) and 2-port 10GB modules

Table 2: Software Branches

Software Release	Includes:	Operates on:
07.7.01b release	Bug fixes, 07.7.01 features and enhancements to existing features, and a new procedure for upgrading software code	ProCurve 9304M, ProCurve 9308M, and ProCurve 9315M routing switches with redundant management (M2, M4, EP, and T-Flow) and 2-port 10GB modules
07.8.00a release	Bug fixes, new 07.8.00a hardware and software features, and enhancements to existing features	ProCurve 9304M, ProCurve 9308M, and ProCurve 9315M routing switches with redundant management (M2, M4, EP, and T-Flow) and 2-port 10GB modules
08.0.01c release	Bug fixes, new features, and enhancements to existing features	ProCurve 9304M, ProCurve 9308M, and ProCurve 9315M routing switches with redundant management (M2, M4, EP, and T-Flow) and 2-port 10GB modules
08.0.01k release	Bug fixes	ProCurve 9304M, ProCurve 9308M, and ProCurve 9315M routing switches with redundant management (M2, M4, EP, and T-Flow) and 2-port 10GB modules

Software Requirements for Management Modules

Table 3 shows the minimum software releases required to run redundant management modules.

Table 3: Minimum Software Requirements for Management Modules

Minimum Software Release Required	ProCurve 9300 Series Redundant Management Modules Supported
07.1.10 Supported only on ProCurve 9304M and 9308M	J4845A ProCurve 9300 GigLX Redundant Management Module (8-port, M2) J4846A ProCurve 9300 GigSX Redundant Management Module (8-port, M2) J4847A ProCurve 9300 Redundant Management Module (0-port, M2)
07.1.19 Supported only on ProCurve 9304M and 9308M	All of the redundant management modules supported for release 07.1.10 J4857A ProCurve 9300 Mini-GBIC Redundant Management Module (8-port, M4)
07.5.04 Supported on ProCurve 9304M, 9308M, and 9315M	All of the redundant management modules supported for release 07.1.19 J4879A ProCurve 9300 T-Flow Redundant Management Module
07.6.00 Supported on ProCurve 9304M, 9308M, and 9315M	All of the redundant management modules supported for release 07.5.04 J4885A ProCurve 9300 EP Mini-GBIC Redundant Management Module

Software Supported on ProCurve Routing Switches

Table 4 shows the software releases supported on each ProCurve routing switch.

Table 4: Software Releases Supported on ProCurve Routing Switches

Routing Switch	Software Releases Supported			
	EP Redundant Mgmt	M2 or M4 Redundant Mgmt	M2 or M4 Redundant Mgmt	M1 Router Code only
	H2R08001c.BIN H2R08001k.BIN	H2R05216.BIN H2R06605.BIN H2R06616.BIN H2R07110.BIN H2R07119.BIN H2R07122.BIN H2R07124.BIN	H2R07504.BIN ¹ H2R07600.BIN ² H2R07601.BIN H2R07604c.BIN H2R07701b.BIN H2R07800a.BIN H2R07801d.BIN	HPR05216.BIN HPR06605.BIN HPR06616.BIN HPR06628.BIN HPR06633.BIN HPR06636.BIN
ProCurve 9315M (J4874A) Routing Switch with EP or Standard (M2 or M4) Redundant Management Module(s)	Yes	No	Yes	No
ProCurve 9304M (J4139A) and 9308M (J4138A) Routing Switches with EP or Standard (M2 or M4) Redundant Management Module(s)	Yes	Yes	Yes	No
ProCurve 9304M (J4139A) and 9308M (J4138A) Routing Switches without Redundant Management (with M1 Management Module)	N/A ³	No	No	Yes
ProCurve 6308M-SX (J4840A) Routing Switch	N/A ³	No	N/A	Yes
ProCurve 6208M-SX (J4841A) Switch	N/A ³	No	N/A	No
¹ First software release to support the ProCurve 9315M routing switch and the J4879A T-Flow module ² First software release to support the EP (Enhanced Performance) modules. ³ Not applicable.				

Boot Code Requirements for ProCurve Software

To run a software release on a ProCurve routing switch, you must use the minimum version of boot code for each software image described in Table 5.

Table 5: Boot Code Requirements

Routing Switch	Modules	Software Image	Minimum Boot Code Required
ProCurve 9304M ProCurve 9308M	With one of the following M1 ¹ modules (without Redundant Management): J4141A 10/100 J4144A Gigabit SX J4146A Gigabit 4LX/4SX	HPR06636.bin ²	M1B07108.bin or greater recommended
ProCurve 9304M ProCurve 9308M ProCurve 9315M	With any one or two of the following Redundant Management: modules J4846A Gigabit SX ² (M2) J4845A Gigabit LX ² (M2) J4847A 0-Port ² (M2) J4857A Mini-GBIC (M4) J4885A EP J4879A T-Flow	For M2, M4, and EP Redundant Management modules: - H2R08001k.bin For T-Flow Redundant Management modules: - TSP08001k.bin	For M2, M4, and EP Redundant Management modules: - M2B07605.bin or greater For T-Flow Redundant Management modules: - M2B07605.bin (all MP images) - VSB07100.bin (VSM code)
—	10 Gigabit Ethernet (10GE) Modules Note: To upgrade FPGA code, refer to “Upgrading the FPGA on a 10 Gigabit Ethernet Module” on page 11.	The Field-Programmable Gate Arrays (FPGAs) in 10 Gigabit Ethernet Modules use the following software. J4891A FPGA: rxmgr.bin – version 80, revision 6 rxpp.bin – version 81, revision 16 txaccum.bin – version 82, revision 6 txpp.bin – version 83, revision 13 ageram.bin – version 84, revision 4 J8174A FPGA: xpp.bin – version 88, revision 42 xpm.bin – version 89, revision 43 Note: To determine the FPGA versions running on a 10GE module, enter show flash . The version information is listed separately for each 10 Gigabit Ethernet module in the chassis.	10 Gigabit Ethernet modules do not use any of the boot images listed above.
ProCurve 6308M-SX	—	HPR06636.bin ²	M1B07108.bin or greater recommended

Routing Switch	Modules	Software Image	Minimum Boot Code Required
ProCurve 6208M-SX	—	HPS06636.bin ²	M1B07108.bin or greater recommended
¹ M1 management modules (without Redundant Management) have been discontinued. ² Does not support Secure Shell (SSH) version 1.			

Upgrading Software on an M2, M4, or EP Management Module to Release 08.0.01k

This section explains how to upgrade the software used on M2, M4, and EP redundant management modules on a ProCurve 9304M, ProCurve 9308M, or ProCurve 9315M routing switch to release 08.0.01k.

NOTE: As shown in Table 5, newer software versions require newer versions of boot code. Software versions use a five-digit number in the format: xx.x.xx; for example, 08.0.01k. Boot code versions use a six-digit number in the format: xx.xx.xx; for example, 07.06.05.

Different procedures are used to upgrade an M2, M4, or EP management module, depending on the version of software running on the module:

- A software release earlier than 07.6.01b
- Software release 07.6.01b or greater

NOTE: M1 Management modules (discontinued) do not support software releases 07.x.xx, and are, therefore, not described in this section. The latest software release supported on an M1 management module is 06.6.36.

Restrictions

- Software release 08.0.01k requires boot code version 07.06.05 to support all hardware modules and decompress new software images.

A new compression algorithm was introduced in software releases greater than 07.6.01b to generate software images. The new compression algorithm allows a software image to contain more features.

Software release 07.6.01b was introduced as a special release that is used as an intermediate step when you upgrade to a later software release. After you install release 07.6.01b and reboot a routing switch, the switch is able to copy the latest software images to flash memory.

- On an M2, M4, or EP redundant management module, boot code is not automatically copied from the active to the standby management module. (However, software code is automatically copied to a standby management module.)

To copy boot code from the active to a standby management module, you must enter the **sync boot** command.

- On a ProCurve 9315M, software release 07.5.04 is the earliest release supported. If a management module is running software earlier than release 07.5.04, you cannot upgrade the module in a 9315M chassis. Instead, you must upgrade it in a 9304M or 9308M chassis.

Upgrading to Software Release 08.0.01k: Overview

To upgrade an M2, M4, or EP management module to release 08.0.01k, you must follow these general steps:

- A.** Upgrade the boot code on the management module to version 07.06.05. If necessary, use the **sync boot** command to copy boot code from the active to a standby management module in the routing switch. Then reboot the routing switch to load boot code 07.06.05.
- B.** If the routing switch is running software EARLIER than release 07.6.01b, copy release 07.6.01b to flash memory. Then reboot the device to load the 07.6.01b software.
- C.** Copy release 08.0.01k to flash memory, and reboot the routing switch to load 08.0.01c software.

A. Upgrading Boot Code on a Management Module to Version 07.06.05

To upgrade the boot code on an M2, M4, or EP management module to version 07.06.05:

1. Store boot code version 07.06.05 (filename: M2B07605.bin) on a TFTP server that the routing switch can access.
2. Enter the following command at the privileged EXEC level of the CLI (for example: ProCurve 9300#) to copy the boot code from the TFTP server into the flash memory of the management module:

```
copy tftp flash <ip-addr> <image-file-name> boot
```

3. Verify that the code has been successfully copied by entering the following command at any level of the CLI:
show flash

The boot code version is displayed on the line that begins with "Boot Image size". Ensure that boot code version 07.06.05 is displayed for the active management module.

4. If a standby (redundant) management module is installed in the routing switch, synchronize the boot code on the standby management module by entering the **sync boot** command.

Verify that boot code 07.06.05 has been successfully copied on the standby management module by entering the **show flash** command.
5. Reboot the routing switch to load boot code 07.06.05.

B. Upgrading Software on a Management Module From a Release Earlier than 07.6.01b

To upgrade the software on an M2, M4, or EP management module from a release EARLIER than 07.6.01b to release 08.0.01k:

1. Verify the version of boot code installed on the management module by entering the **show flash** command.

The boot code version is displayed at the end of the line that begins with "Boot Image size". Ensure that boot code version 07.06.05 is displayed.

NOTE: The **show flash** command only displays the version of boot code installed on the device. It does not display the version of boot code running on the device.

If you rebooted the routing switch after installing boot code 07.06.05 (as described in "A. Upgrading Boot Code on a Management Module to Version 07.06.05"), the required boot code is running. If you are not sure, ProCurve recommends that you reboot the device now.

2. Store software release 07.6.01b (filename: H2R07601b.bin) on a TFTP server that the routing switch can access.

3. Upgrade the software on the management module to version 07.6.01b by entering the following command:

```
copy tftp flash <ip-address> H2R07601b.bin [primary | secondary]
```

Where:

primary copies software to the primary (default) storage area in flash memory.

secondary copies software to the secondary area in flash memory.

If no redundant management module is installed, the message `TFTP to Flash Done` is displayed when the upgrade is complete.

If a redundant management module is installed, the message `Sync Secondary code in flash...Done` is displayed when the flash images are synchronized and the upgrade is complete.

4. Verify that the software has been successfully copied by entering the **show flash** command at any level of the CLI:
 - The software release in the primary flash is displayed at the end of the line that begins with "Compressed Pri Code Size".
 - The software release in the secondary flash is displayed at the end of the line that begins with "Compressed Sec Code Size".

Ensure that software release 07.6.01b is stored in the primary or secondary flash area.

5. Reboot the routing switch to load software release 07.6.01b from the area of flash memory (primary or secondary) where you stored it.
6. Continue with "C. Upgrading Software on a Management Module From Release 07.6.01b or Greater" to upgrade the software to release 08.0.01k.

C. Upgrading Software on a Management Module From Release 07.6.01b or Greater

To upgrade the software on an M2, M4, or EP management module from release 07.6.01b or greater to release 08.0.01k:

1. Verify the version of boot code running on the management module by entering the **show flash** command.

The boot code version is displayed at the end of the line that begins with "Boot Image size". Ensure that boot code version 07.06.05 is displayed.

NOTE: The **show flash** command only displays the version of boot code installed on the device. It does not display the version of boot code running on the device.

If you rebooted the routing switch after installing boot code 07.06.05 (as described in "A. Upgrading Boot Code on a Management Module to Version 07.06.05"), the required boot code is running. If you are not sure, ProCurve recommends that you reboot the device now.

2. Store software release 08.0.01k (filename: H2R08001c.bin) on a TFTP server that the routing switch can access.
3. Upgrade the software on the management module to version 08.0.01k by entering the following command:

```
copy tftp flash <ip-address> H2R08001c.bin [primary | secondary]
```

Where:

primary copies software to the primary (default) storage area in flash memory.

secondary copies software to the secondary storage area.

If no redundant management module is installed, the message `TFTP to Flash Done` is displayed when the upgrade is complete.

If a redundant management module is installed, the message `Sync Secondary code in flash...Done` is displayed when the flash images are synchronized and the upgrade is complete.

4. Verify that the software has been successfully copied by entering the **show flash** command at any level of the CLI:
 - The software release in the primary flash is displayed at the end of the line that begins with “Compressed Pri Code Size”.
 - The software release in the secondary flash is displayed at the end of the line that begins with “Compressed Sec Code Size”.

Ensure that software release 08.0.01k is stored in the primary or secondary flash area.

5. Reboot the routing switch to load software release 08.0.01k from the area of flash memory (primary or secondary) where you stored it.

NOTE: When you reload the software after upgrading the software to release 08.0.01k, the routing switch displays a message to say that the configuration has changed and prompts you to save the changes. This message is displayed even if you do not make any configuration changes because the software records its release number in the running-config file when the software is loaded onto the switch. Enter **Y** to reload without saving the change or save the change and reload.

Using SNMP to Upgrade Software on a Management Module

Third-party SNMP management applications such as HP OpenView can upgrade software on a routing switch.

NOTE: In software releases earlier than 07.5.04, the SNMP agent does not check for type validity with the SNMP version. In software release 07.5.04 and greater, the SNMP agent does not send a reply for a varbind, if the type of the varbind is not a known type for that version of SNMP. For example, MIB objects of type Counter64 cannot be retrieved using a v1 packet, as Counter64 is a v2c and v3 type.

Make sure you use the correct procedure for your device and processor type. For example, do not use the Management Processor procedure to upgrade the switching processors on a T-Flow module.

The syntax shown in this section assumes that you have installed HP OpenView in the “/usr” directory.

ProCurve recommends that you make a backup copy of the startup-config file before you upgrade the software. If you need to run an older release, you will need to use the backup copy of the startup-config file.

Upgrading a Management Processor using SNMP

Use the following procedure to upgrade:

- An M2, M4, or EP module
- Management processor on the T-Flow module

To upgrade software code on the Management Processor:

1. Configure a read-write community string on the ProCurve device, if one is not already configured. To configure a read-write community string, enter the following command from the global CONFIG level of the CLI:

```
snmp-server community <string> rw
```

where *<string>* is the community string and can be up to 32 characters long.

2. On the ProCurve device, enter the following command from the global CONFIG level of the CLI:

```
no snmp-server pw-check
```

This command disables password checking for SNMP set requests. If a third-party SNMP management application does not add a password to the password field when it sends SNMP set requests to an ProCurve device, by default the ProCurve device rejects the request.

- From the command prompt in the UNIX shell, enter the following command:

```
/usr/OV/bin/snmpset -c <rw-community-string> <hp-ip-addr> 1.3.6.1.4.1.1991.1.1.2.1.5.0  
ipaddress <tftp-ip-addr> 1.3.6.1.4.1.1991.1.1.2.1.6.0 octetstringascii <file-name>  
1.3.6.1.4.1.1991.1.1.2.1.7.0 integer <command-integer>
```

Where:

<rw-community-string> is a read-write community string configured on the ProCurve device.

<hp-ip-addr> is the ProCurve device's IP address.

<tftp-ip-addr> is the TFTP server's IP address.

<file-name> is the image file name.

<command-integer> is one of the following values:

20 – Downloads the software code into the device's primary flash area.

22 – Downloads the software code into the device's secondary flash area.

Upgrading the FPGA on a 10 Gigabit Ethernet Module

This section explains how to upgrade an FPGA (Field-Programmable Gate Array) on a 10 Gigabit Ethernet module. 10 Gigabit Ethernet modules do not have boot code separate from the management module. However, they do have FPGAs that require separate software.

NOTE: The J8174A 2-port 10 Gigabit Ethernet module with XENPAK optics uses a different FPGA file than the older J4891A 1-port 10 Gigabit Ethernet module. See Table 5 on page 6 for a list of the FPGA files supported on both the 1-port and 2-port 10 Gigabit Ethernet modules.

The J8174A 2-port 10 Gigabit Ethernet module with XENPAK optics can function in the same chassis with the older J4891A 1-port 10 Gigabit Ethernet module.

If an upgrade is required for any of the FPGA files, you must upgrade all the FPGA files.

To upgrade the FPGA on a 10 Gigabit Ethernet module:

- Complete the upgrades of the boot code and software code, if required.
- Enter the following command for the 10 Gigabit Ethernet module at the privileged EXEC level of the CLI:

```
10gig copy tftp flash <ip-addr> <filename> [module <slotnum>]
```

Where:

tftp specifies the location of the FPGA file. The **tftp** parameter shows that the file is on a TFTP server.

<ip-addr> specifies the IP address of the TFTP server, if you specify **tftp**.

<filename> specifies the FPGA file name. The 2-port 10 Gigabit Ethernet module has only two FPGA files; xpp.bin and xtm.bin. For more information on the supported FPGA files, see Table 5 on page 6.

module <slotnum> is an optional parameter that specifies the modules on which you want to install the upgrade. If you do not specify a slot number, the command upgrades the FPGA on all 10 Gigabit Ethernet modules in the chassis.

Example

```
ProCurve 9300# 10gig copy tftp flash 10.10.10.10 rxbmgr.bin
ProCurve 9300# 10gig copy tftp flash 10.10.10.10 rxpp.bin
ProCurve 9300# 10gig copy tftp flash 10.10.10.10 txaccum.bin
ProCurve 9300# 10gig copy tftp flash 10.10.10.10 txpp.bin
ProCurve 9300# 10gig copy tftp flash 10.10.10.10 ageram.bin
```

NOTE: You can store and copy the FPGA files using any valid filename. You are not required to store and copy the files using the names listed in “Boot Code Requirements for ProCurve Software” on page 5. The device uses information within the files to install them in the correct FPGAs, and the **show flash** command lists the FPGAs according to the names in “Boot Code Requirements for ProCurve Software” on page 5.

3. Reload the software by entering one of the following commands:
 - **reload** (this command boots from the default boot source, which is the primary flash area by default)
 - **boot system flash primary | secondary**
-

NOTE: The **show flash** command will list the new FPGA code versions but the new versions do not take effect until you reload the software.

Using Different Combinations of Management Modules

This section describes the different combinations of M1, M2, M4, EP, and T-Flow management modules supported on ProCurve 9304M, ProCurve 9308M, and ProCurve 9315M routing switches.

Table 6: Supported Management Module Combinations

Primary Management Module	Secondary Management Module	Notes
J4885A EP Mini-GBIC Redundant Management Module	Another EP Redundant Management Module	—
Any M2 or M4 Redundant Management Module (<i>Discontinued</i>)	Another M2 or M4 Redundant Management Module	When you use an M2 and M4 in the same switch, ProCurve recommends using the faster M4 as the primary redundant management module. If the M4 fails, the system will use the slower M2 module.
J4879A T-Flow Redundant Management Module (<i>Discontinued</i>)	Another J4879A T-Flow Redundant Management Module	—
Any M1 Management Module (<i>Discontinued</i>)	N/A	Supported only in the ProCurve 9304M and ProCurve 9308M routing switches. No redundant management options.

NOTE: The following types of management modules are *mutually exclusive*:

- M1 management modules
- M2 or M4 redundant management modules
- EP redundant management modules
- T-Flow redundant management modules

A ProCurve routing switch does not operate if two redundant management modules of different types are installed. Also, M1 management modules do not operate in a ProCurve 9315M routing switch.

Redundant management means that a device can operate with two management modules installed; one active (primary) and one standby (secondary). If the active management module becomes unavailable, the standby management module automatically takes over system operation.

Management modules with redundant management capabilities include the following M2, M4, EP, and T-Flow modules:

- J4885A ProCurve 9300 EP Mini-GBIC Redundant Management Module (8-port)
- J4879A ProCurve 9300 T-Flow Redundant Management Module (0-port — *discontinued*)
- J4857A ProCurve 9300 Mini-GBIC Redundant Management Module (8-port, M4 — *discontinued*)
- J4845A ProCurve 9300 GigLX Redundant Management Module (8-port, M2 — *discontinued*)
- J4846A ProCurve 9300 GigSX Redundant Management Module (8-port, M2 — *discontinued*)
- J4847A ProCurve 9300 Redundant Management Module (0-port, M2 — *discontinued*)

If you are using a Redundant Management module, you can install either one or two such modules in the routing switch, as shown in Table 6. For more information, see “Using Redundant Management Modules” in the *Installation and Basic Configuration Guide* included on the *Documentation CD-ROM* shipped with your management module, and also downloadable from the ProCurve website (see “To Download Product Documentation:” on page 3).

Non-Redundant Management on ProCurve 9304M and ProCurve 9308M Routing Switches

Management modules without Redundant Management are sometimes termed “M1” modules (for “Management 1”). These modules, now discontinued, operate only in the ProCurve 9304M and ProCurve 9308M routing switches. M1 modules include:

- J4141A ProCurve 9300 10/100 Management Module (16-port — *discontinued*)
- J4144A ProCurve 9300 Gigabit SX Management Module (8-port — *discontinued*)
- J4146A ProCurve 9300 Gigabit 4LX/4SX Management Module (8-port — *discontinued*)

NOTE: M1 management modules do not operate in the ProCurve 9315M routing switch. Also, if you are using an M1 management module in a ProCurve 9304M or ProCurve 9308M, no other management module (non-redundant or redundant) can be installed in the routing switch.

Maximum Size of Startup-Config and Running-Config Files

Each ProCurve device has a maximum supported size for the running-config and the startup-config file. If you use TFTP to load additional information into a device’s running-config or startup-config file, it is possible to exceed the maximum supported size. If this occurs, you will not be able to save the configuration changes.

Table 7 lists the maximum size of the running-config and the startup-config files on ProCurve devices.

Table 7: Maximum Sizes Supported for running-config and the startup-config Files

Device	Maximum Size of running-config and startup-config files
ProCurve 9315 using Management II or higher	256 kilobytes (KB)
ProCurve 9304M or ProCurve 9308M using Management II or higher	256 KB
ProCurve 9304M or ProCurve 9308M using Management I (<i>discontinued</i>)	128 KB
ProCurve 6308M-SX or ProCurve 6208M-SX (<i>discontinued</i>)	64 KB

NOTE: The maximum supported file size of each file is not the combined size of the running-config and startup-config files. The running-config and startup-config files can each be the size listed in Table 7.

To determine the size of an ProCurve device's running-config or startup-config file, copy the file to a TFTP server. Then use the directory services on the server to list the size of the copied file. To copy the running-config or startup-config file to a TFTP server, use one of the following commands.

- To copy the running-config to a TFTP server:
copy running-config tftp <ip-addr> <filename>
- To copy the startup-config file to a TFTP server:
copy startup-config tftp <ip-addr> <filename>

Removing a Module from an Active Device

Before you remove a module from a routing switch in operation, first disable the module. Disabling the module before removing it prevents a brief service interruption on other unmanaged modules. The brief interruption can be caused by the device re-initializing other modules when you remove an enabled module.

NOTE: This section does not apply to the active or standby Redundant Management modules. The **disable module** and **enable module** commands are not supported on management modules.

To disable a module, enter the following command at the Privileged EXEC level of the CLI:

disable module <slot-num>

Where <slot-num> specifies the slot number as follows:

- Slots in a 4-slot chassis are numbered 1 – 4, from top to bottom.
- Slots in an 8-slot chassis are numbered 1 – 8, from left to right.
- Slots in a 15-slot chassis are numbered 1 – 15, from left to right.

Example

```
ProCurve 9300# disable module 3
```

This command disables the module in slot 3.

NOTE: If you remove the module without first disabling it, the routing switch re-initializes the other modules in the device, causing a brief interruption in service after which the device resumes normal operation.

You do not have to enable a module after inserting it in the chassis. The module is automatically enabled when you insert the module into a running device or when you power on the device.

To replace a removed module with a different type of module, you must configure the slot for the new module. To configure a slot for a module, use the **module** command at the global CONFIG level of the CLI.

After disabling a module, if you decide not to remove the module, re-enable the module using the following command:

enable module <slot-num>

Example

The following command re-enables the module in slot 3:

```
ProCurve 9300# enable module 3
```

Configuring the ProCurve 9315M

When configuring a ProCurve 9315M 15-slot routing switch, take into account the guidelines and restrictions in this section.

Minimum Software Release Supported

The ProCurve 9315M requires software release 07.5.04 or greater.

NOTE: On a ProCurve 9315M, software release 07.5.04 is the earliest release supported. If a management module is running software earlier than release 07.5.04, you cannot upgrade the module in a 9315M chassis. Instead, you must upgrade it in a 9304M or 9308M chassis. For more information, see “Upgrading Software on an M2, M4, or EP Management Module to Release 08.0.01k” on page 7.

Inserting or Removing an EP Module on a ProCurve 9315M

NOTE: This section applies only to a ProCurve 9315M (15-slot chassis) with EP modules.

Do not insert or remove EP modules in a ProCurve 9315M until the device has fully booted. Generally, booting takes around two minutes. You can determine whether the device has fully booted by looking at the management console. Once the device boots, a command prompt or login prompt is displayed.

After the device has booted, allow time for the device to fully complete the removal or insertion before removing or inserting another module. Generally, this takes about 30 seconds. After you remove or insert a module, the CLI displays a message confirming completion of the change. Wait for this message before removing or inserting another module.

Slot Locations for Redundant Management Modules

The 15 slots in the ProCurve 9315M are divided among 4 internal regions. Slots 1 – 4 belong to the same region; slots 5 – 8 belong to the same region; slots 9 – 12 belong to the same region, and slots 13 – 15 belong to the same region. If you are using redundant management modules, ProCurve recommends that you place both management modules in slots belonging to the same region. For example, if you place one management module in slot 5, ProCurve recommends that you place the other management module in slot 6, 7, or 8.

MAC Addresses

The ProCurve 9315M makes use of locally administered MAC addresses. If your site already uses locally administered MAC addresses of the vendor OUI, which is 00e052, there could be a MAC address conflict with one of the ports on the ProCurve device.

Server Trunk Groups

If you plan to configure ports on a module into a server trunk group, use the following guideline:

- For a multi-slot trunk group (one configured on two forwarding modules), the modules must both be in the same set of slots (slots 1 – 7 or 9 – 15).

You do not need to follow this guideline for a switch trunk group.

NOTE: In software releases earlier than 07.6.00, the management module(s) and the module that had the server trunk group's ports were required to be in the same set of slots (slots 1 – 7 or 9 – 15). In software release 07.6.00 and later, there is no longer a restriction on the location of the management module relative to the module used for server trunking. However, it is still a requirement that the module that has the server trunk group's lead ports cannot reside in slot 8.

VLANs

In release 07.6.01b, you could configure only up to 2195 Layer 2 VLANs on the ProCurve 9315M routing switch. The **system-max vlan <num>** command allowed you to allocate a higher number of VLANs, but the software allowed you to actually create only 2195 of the allocated VLANs.

Starting with release 07.6.04, this restriction has been removed. You can create the full number of allocated Layer 2 VLANs on the 15-slot ProCurve 9315M routing switch, up to 4095.

Summary of Enhancements in 08.0.00

Layer 2 Enhancements in 08.0.00

Enhancement	Description	EP	Non-EP	See Page
Dynamic configuration of Cisco Voice over IP (VoIP) phones via voice VLANs	This release enables the ProCurve device to automatically re-configure a VoIP phone when the phone is physically moved from one port to another.	✓	✓	26
EP Layer 2 Access Control Lists (ACLs)	You can configure Layer 2 ACLs, which filter incoming traffic based on Layer 2 MAC header fields in the Ethernet/IEEE 802.3 frame, including source MAC address and mask, destination MAC address and mask, VLAN ID, and Ethernet type.	✓		27
Layer 2 ACL-based rate limiting	This feature is an extension to the existing IP ACL-based rate limiting on EP devices. Whereas the existing feature provides the facility to limit the rate for IP traffic that matches the permit conditions in standard or extended IP ACLs; the new feature enables you to limit traffic rates using the Layer 2 parameters defined in the associated EP Layer 2 ACL table.	✓		30
802.1s support	Multiple Spanning Tree Protocol (MSTP) as defined in IEEE 802.1s-2002 allows you to configure multiple STP instances.	✓	✓	32

Layer 3 Enhancements in 08.0.00

Enhancement	Description	EP	Non-EP	See Page
Clearing OSPF Information from the ProCurve Device	This release includes new CLI commands that allow you to clear specific kinds of information from the ProCurve device's OSPF link state database and OSPF routing table. You do not need to remove statements from the ProCurve device's configuration or reload the software for the commands to take effect.	✓	✓	19

Enhancement	Description	EP	Non-EP	See Page
OSPF redistribution filter rebinding	In previous releases, if you modified OSPF redistribution filters in the ProCurve device's configuration, you then had to remove and reapply the redistribution rip and redistribution static statements in order for OSPF to start redistributing routes based on the filters. Starting in release 08.0.00, a new CLI command rebinds the redistribution filters, so you no longer have to remove and reapply the redistribution rip and redistribution static statements.	✓	✓	21
Configuring an OSPF non-broadcast interface	You can configure an interface on a ProCurve device to send OSPF traffic to its neighbor as unicast packets rather than broadcast packets.	✓	✓	21
Configuring a unicast route with multiple outgoing ports	In this release, the ProCurve device can route incoming Layer 3 unicast IP packets to two or more statically defined outgoing Layer 3 interfaces. Previous releases support sending traffic received from a VLAN or a Layer 3 routed interface to multiple Layer 2 addresses.	✓	✓	22
Advertising an IBGP next hop as a null0 route as a defense against denial of service attacks	You can create a static route to forward traffic destined to a host or network targeted in a denial of service attack to the null0 interface, and then use IBGP to advertise the null0 route to other routers in the network, so that the other routers drop the traffic for the targeted host or network.	✓	✓	23
IP load sharing (ECMP) for RIPv2 routes	In release 08.0.00, IP load sharing is supported for RIPv2 routes. IP load sharing is also known as "Equal-Cost Multi-Path (ECMP)" load sharing or just "ECMP".	✓	✓	25
VRRPE slow start timer	The VRRPE slow start timer causes a specified amount of time to elapse between the time a VRRPE Master router comes back up and when it takes over from a Backup router. This interval allows time for OSPF convergence when the Master is restored.	✓	✓	25

System-Level Enhancements in 08.0.00

Enhancement	Description	EP	Non-EP	See Page
Creating an alias for a CLI command	You can create aliases for CLI commands. An alias serves as a shorthand version of a longer CLI command; for example, shoro can be defined as an alias for the show ip route CLI command.	✓	✓	41

Enhancement	Description	EP	Non-EP	See Page
Directing debugging output to multiple destinations	You can direct debugging output (output from debug commands) to multiple destinations. This allows debugging output to be displayed on multiple user sessions concurrently. In previous releases, debugging output could be directed only to a single destination (the Syslog buffer, or a specified Telnet or SSH session).	✓	✓	42
sFlow enhancements	Release 08.0.00 includes the following enhancements to ProCurve's support for sFlow: <ul style="list-style-type: none"> Support for sFlow version 5 Support for selecting which and how often data destined to the CPU (for example, Telnet sessions) is sent to the sFlow collector 	✓		43
Jumbo packet counter in show rmon statistics output	The output of the show rmon statistics command has been enhanced to display information about the number of packets received that were longer than 1518 octets.	✓		45
Displaying VLANs in numerical order	The output of the show run command displays the VLANs configured on the ProCurve device in numerical order (for example, VLAN 100, VLAN 200, and so on).	✓	✓	46
Specifying a host name in an ACL statement	You can specify a host name within the configuration of an access control list. Previous releases supported only IP addresses in ACL statements.	✓	✓	46
Multiple domain list support	You can create a list of domain names that can be used by the DNS Resolver to resolve host names with their IP addresses. You can also verify the host name for an IP address or the IP address for a host name.	✓	✓	47

Multicast Enhancements in 08.0.00

Enhancement	Description	EP	Non-EP	See Page
IGMP V3 Snooping	Release 08.0.00 supports IGMP V3 snooping.	✓		54
Increased size of the MSDP Source Active cache	The size of the cache used to store MSDP Source Active messages has been increased from 4k to 8k.	✓	✓	63
Specifying a Designated Router election priority for PIM	You can assign a DR election priority to each PIM router in a multi-access network. The router with the highest DR election priority is elected the DR.	✓	✓	63

Enhancement	Description	EP	Non-EP	See Page
Multicast MIB support	Support for the following standard multicast MIBs has been added to this release: <ul style="list-style-type: none"> • RFC 2932: IPv4 Multicast Routing MIB • RFC 2933: IGMP MIB • RFC 2934: PIM for IPv4 	✓	✓	64

Security Enhancements in 08.0.00

Enhancement	Description	EP	Non-EP	See Page
Using multi-device port authentication and 802.1X security on the same port	You can configure the ProCurve device to use multi-device port authentication and 802.1X security on the same port.	✓	✓	69
CPU ACL	Starting in release 08.0.00, you can configure the ProCurve device to protect itself from denial of service attacks for all protocols and applications received on the device. To do so, extended ACLs applied to ingress ports filter on the protocols specified in ACL clauses. For example, you can use extended ACLs to filter on protocols such as SNMP, TFTP, BGP, and ICMP.	✓		75

Enhancements and Configuration Notes in 08.0.00

Clearing OSPF Information from the ProCurve Device

Releases 08.0.00 and later include new CLI commands that allow you to clear specific kinds of information from the ProCurve device's OSPF link state database and OSPF routing table.

The following kinds of OSPF information can be cleared:

- Routes received from OSPF neighbors. You can clear routes from all OSPF neighbors, or an individual OSPF neighbor, specified either by the neighbor's IP address or its router ID
- OSPF topology information, including all routes in the OSPF routing table
- All routes in the OSPF routing table that were redistributed from other protocols
- OSPF area information, including routes received from OSPF neighbors within an area, as well as routes imported into the area. You can clear area information for all OSPF areas, or for a specified OSPF area

The OSPF information is cleared dynamically when you enter the command; you do not need to remove statements from the ProCurve device's configuration or reload the software for the change to take effect.

Clearing OSPF Neighbor Information

To clear information on the ProCurve device about all OSPF neighbors, enter the following command:

```
ProCurve 9300# clear ip ospf neighbor
```

Syntax: clear ip ospf neighbor [ip <ip-addr> | id <ip-addr>] |

This command clears all OSPF neighbors and the OSPF routes exchanged with the neighbors in the ProCurve device's OSPF link state database. After this information is cleared, adjacencies with all neighbors are re-established, and routes with these neighbors exchanged again.

To clear information on the ProCurve device about OSPF neighbor 10.10.10.1, enter the following command:

```
ProCurve 9300# clear ip ospf neighbor ip 10.10.10.1
```

This command clears the OSPF neighbor and the OSPF routes exchanged with neighbor 10.10.10.1 in the ProCurve device's OSPF link state database. After this information is cleared, the adjacency with the neighbor is re-established, and routes are exchanged again.

The neighbor router can be specified either by its IP address or its router ID. To specify the neighbor router using its IP address, use the **ip** <ip-addr> parameter. To specify the neighbor router using its router ID, use the **id** <ip-addr> parameter.

Clearing OSPF Topology Information

To clear OSPF topology information on the ProCurve device, enter the following command:

```
ProCurve 9300# clear ip ospf topology
```

Syntax: clear ip ospf topology

This command clears all OSPF routes from the OSPF routing table, including intra-area, (which includes ABR and ASBR intra-area routes), inter-area, external type 1, external type 2, OSPF default, and OSPF summary routes.

After you enter this command, the OSPF routing table is rebuilt, and valid routes are recomputed from the OSPF link state database. When the OSPF routing table is cleared, OSPF routes in the global routing table are also recalculated. If redistribution is enabled, the routes are imported again.

Clearing Redistributed Routes from the OSPF Routing Table

To clear all routes in the OSPF routing table that were redistributed from other protocols, enter the following command:

```
ProCurve 9300# clear ospf redistribution
```

Syntax: clear ospf redistribution

This command clears all routes in the OSPF routing table that are redistributed from other protocols, including direct connected, static, RIP, and BGP. To import redistributed routes from other protocols, use the **redistribution** command at the OSPF configuration level.

Clearing Information for OSPF Areas

To clear information on the ProCurve device about all OSPF areas, enter the following command:

```
ProCurve 9300# clear ip ospf
```

This command clears all OSPF areas, all OSPF neighbors, and the entire OSPF routing table. After this information has been cleared, adjacencies with all neighbors are re-established, and all OSPF routes are re-learned.

To clear information on the ProCurve device about OSPF area 1, enter the following command:

```
ProCurve 9300# clear ip ospf area 1
```

This command clears information about the specified area ID. Information about other OSPF areas is not affected. The command clears information about all OSPF neighbors belonging to the specified area, as well as all routes imported into the specified area. Adjacencies with neighbors belonging to the area are re-established, and routes imported into the area are re-learned.

Syntax: clear ip ospf area [<area-id>]

The <area-id> can be specified in decimal format or in IP address format.

OSPF Redistribution Filter Rebinding

In previous releases, if you modified OSPF redistribution filters in the ProCurve device's configuration, you then had to remove and reapply the **redistribution rip** and **redistribution static** statements in order for OSPF to start redistributing routes based on the filters.

Starting in release 08.0.00, a new CLI command rebinds the redistribution filters, so you no longer have to remove and reapply the **redistribution rip** and **redistribution static** statements in the ProCurve device's configuration.

For example, if the ProCurve device's configuration contained the following statements:

```
permit redistribute 1 static address 130.126.0.12 255.255.255.255
permit redistribute 2 static address 128.174.201.0 255.255.255.128
permit redistribute 3 rip
deny redistribute 64 all
redistribution rip
redistribution static
```

and you then added the following filter statement:

```
permit redistribute 4 static address 192.17.220.0 255.255.254.0
```

you then had to remove the **redistribution rip** and **redistribution static** statements and then re-enter them in order for OSPF to redistribute routes based on the filter statements.

In this release, instead of removing and reapplying the **redistribution rip** and **redistribution static** statements, you can enter the following commands after modifying OSPF redistribution filters:

```
ProCurve 9300(config)# router ospf
ProCurve 9300(config-ospf-router)# redistribution rebind
```

Syntax: redistribution rebind

After you enter the **redistribution rebind** command, OSPF redistributes routes based on all of the filter statements in the configuration. Note that the **redistribution rebind** command is not stored in the ProCurve device's configuration.

Configuring an OSPF Non-Broadcast Interface

Starting in release 08.0.00, you can configure an interface on a ProCurve device to send OSPF traffic to its neighbor as unicast packets rather than broadcast packets.

OSPF routers generally use broadcast packets to establish neighbor relationships and broadcast route updates on Ethernet and virtual interfaces (VEs). In this release, as an alternative, you can configure the ProCurve device to use unicast packets for this purpose. This can be useful in situations where multicast traffic is not feasible (for example when a firewall does not allow multicast packets).

On a non-broadcast interface, the routers at the other end of this interface must configure non-broadcast and neighbor. There is no restriction on the number of routers sharing a non-broadcast interface (for example, through a hub/switch).

To configure an OSPF interface as a non-broadcast interface, you enable the feature on a physical interface or a VE, following the **ip ospf area** statement, and then specify the IP address of the neighbor in the OSPF configuration. The non-broadcast interface configuration must be done on the OSPF routers on both ends of the link.

For example, the following commands configure VE 20 as a non-broadcast interface:

```
ProCurve 9300(config)# int ve 20
ProCurve 9300(config-vif-20)# ip ospf area 0
ProCurve 9300(config-vif-20)# ip ospf network non-broadcast
ProCurve 9300(config-vif-20)# exit
```

Syntax: [no] ip ospf network non-broadcast

The following commands specify 1.1.20.1 as an OSPF neighbor address. The address specified must be in the same sub-net as a non-broadcast interface.

```
ProCurve 9300(config)# router ospf
ProCurve 9300(config-ospf-router)# neighbor 1.1.20.1
```

For example, to configure the feature in a network with three routers connected by a hub or switch, each router must have the linking interface configured as a non-broadcast interface, and both of the other routers must be specified as neighbors.

The output of the **show ip ospf interface** command has been enhanced to display information about non-broadcast interfaces and neighbors that are configured in the same sub-net.

For example:

```
ProCurve 9300# show ip ospf interface
v20,OSPF enabled
  IP Address 1.1.20.4, Area 0
  OSPF state BD, Pri 1, Cost 1, Options 2, Type non-broadcast Events 6
  Timers(sec): Transit 1, Retrans 5, Hello 10, Dead 40
  DR: Router ID 1.1.13.1 Interface Address 1.1.20.5
  BDR: Router ID 2.2.2.1 Interface Address 1.1.20.4
  Neighbor Count = 1, Adjacent Neighbor Count= 2
  Non-broadcast neighbor config: 1.1.20.1, 1.1.20.2, 1.1.20.3, 1.1.20.5,
  Neighbor: 1.1.20.5
  Authentication-Key:None
  MD5 Authentication: Key None, Key-Id None, Auth-change-wait-time 300
```

In the Type field, “non-broadcast” indicates that this is a non-broadcast interface. When the interface type is non-broadcast, the Non-broadcast neighbor config field displays the neighbors that are configured in the same sub-net. If no neighbors are configured in the same sub-net, a message such as the following is displayed:

```
***Warning! no non-broadcast neighbor config in 1.1.100.1 255.255.255.0
```

Configuring a Unicast Route With Multiple Outgoing Ports

Previous releases support sending traffic received from a VLAN or a Layer 3 routed interface to multiple Layer 2 addresses. In this kind of configuration, if a packet arrives within a VLAN that had a unicast IP destination address, but multicast MAC addresses, you can statically configure the multicast MAC address on two separate ports, and thus enable “flooding” on the two individual interfaces. If the packet arrives from another VLAN or Layer 3 routed interface, then the ProCurve device takes the routed packet and applies the same multicast MAC “flooding” to the two statically defined ports.

Starting in release 08.0.00, this functionality is now supported at Layer 3. The ProCurve device can route incoming Layer 3 unicast IP packets to two or more statically defined outgoing ports. To configure the ProCurve device to do this, you create a static ARP entry that specifies multiple ports for an IP address. The **arp** command has been enhanced to allow you to specify multiple output ports.

Notes

- There is no limitation on the type of MAC address (Layer 2 multicast, unicast) that can be used with this feature, nor is there a limitation on the number of outgoing ports.
- This feature can be used only in a Layer 3 topology; it cannot be used in a mixed Layer 2 and Layer 3 topology. The multiple outgoing ports must be routed interfaces. Sending unicast Layer 2 traffic to multiple outgoing ports (by statically configuring the same MAC address on two or more ports) has been supported in previous releases.
- The multiple outgoing ports must be Ethernet interfaces. The multiple outgoing ports cannot be virtual interfaces (VEs) or trunk groups.
- This feature is not supported for the default route (0.0.0.0/0).

Configuring a Static ARP Entry with Multiple Outgoing Ports

To create a route that sends unicast Layer 3 traffic to multiple outgoing ports, you create a static ARP entry that specifies multiple ports for an IP address. The **arp** command has been enhanced to allow multiple output ports.

For example, to create a static ARP entry that has output ports of 1/21 and 1/22, enter the following command:

```
ProCurve 9300(config)# arp 1 20.20.20.2 0004.809e.2e15 multi-ports e 1/21 e 1/22
```

You can also specify the output ports as a range. For example:

```
ProCurve 9300(config)# arp 1 200.200.200.2 0004.809e.2e15 multi-ports e 1/21 to 1/22
```

Syntax: [no] arp <num> <ip-addr> <mac-addr> ethernet <portnum> | multi-ports ethernet <portnum> [to <portnum>] [ethernet <portnum>]

Displaying Information About Unicast Routes with Multiple Outgoing Ports

You can use the **show arp** and the **show ip route** commands to display information about the unicast routes configured with multiple outgoing ports.

For example, in the following **show arp** output, there are two entries for a single route, both with a different outgoing port:

```
ProCurve 9300# show arp
Total number of ARP entries: 1
   IP Address      MAC Address      Type      Age      Port
1    1.1.21.2       0004.809e.2e15   Static    None     1/21
1    1.1.21.2       0004.809e.2e15   Static    None     1/22
```

In previous releases, the **show arp** command displayed a single line for each ARP entry.

Advertising an IBGP Next Hop as a null0 Route as a Defense Against DDoS Attacks

In a distributed denial of service (DDoS) attack, a substantial amount of traffic may be directed at a targeted host or network. In this situation, you can create a static route to forward traffic destined to the targeted host or network to the null0 interface. Traffic forwarded to the null0 interface is dropped in hardware.

Starting in release 08.0.00, you can also use IBGP to advertise the null0 route to other routers in the network, so that the other routers drop the traffic for the targeted host or network.

To do this, you perform the following tasks:

1. Configure a null0 route on the other routers in the network.
2. On one of the ProCurve routers, configure a static route for the targeted host or network. The static route is configured with a tag so that only specific routes will be redistributed via IBGP, rather than all static routes.
3. Create a route map that matches the tag, sets a local preference value that causes the route to be the preferred route, and configures the origin as IBGP.
4. Redistribute the routes matching the route map via IBGP.

NOTE: This feature applies only to Layer 3 network prefixes. All services for the targeted network prefix are filtered out. To use Layer 4 information as criteria for discarding packets, use an extended ACL.

For example, Figure 1 illustrates a configuration where host 20.20.20.20 is targeted in a DDoS attack.

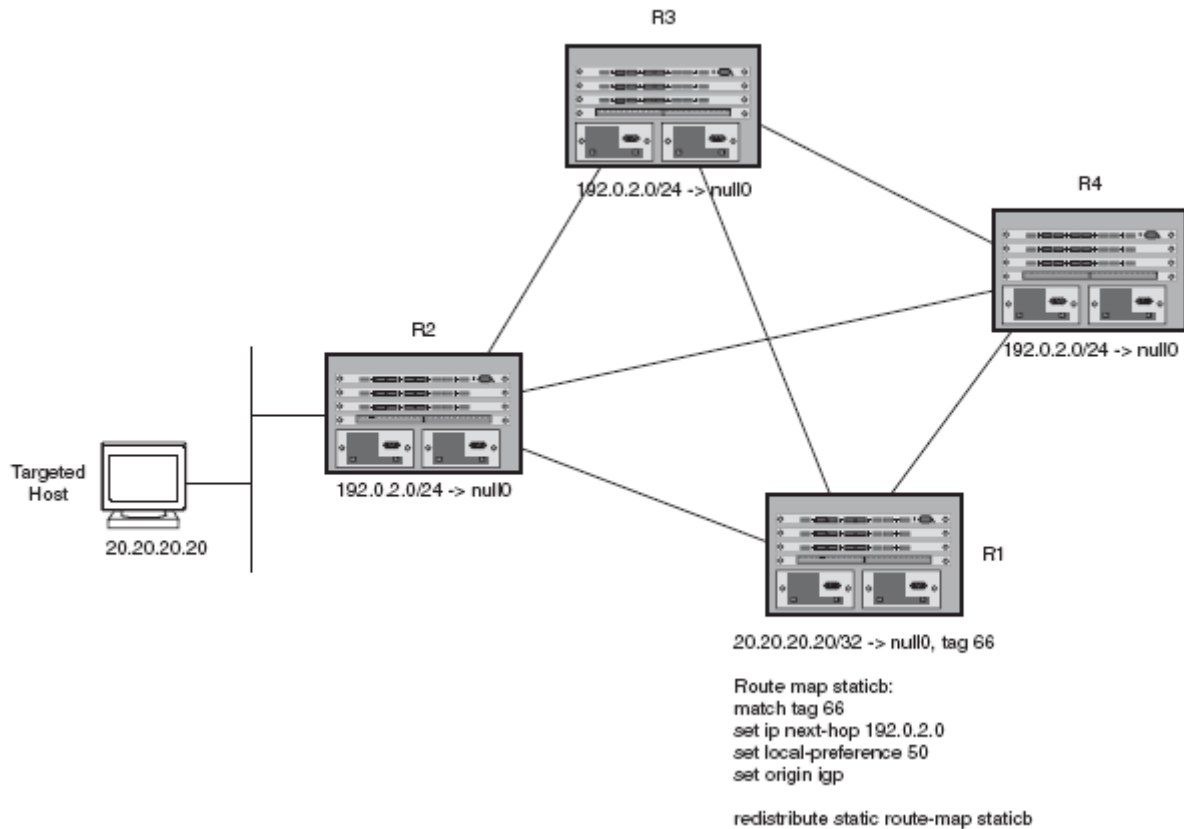


Figure 1 Host Targeted in a DDoS Attack

In this configuration, routers R2, R3, and R4 have a static route that sends packets for 192.0.2.0/24 to the null0 interface. On a ProCurve device, this static route is configured with the following command:

```
ProCurve 9300(config)# ip route 192.0.2.0 255.255.255.0 null0
```

On router R1, a static route is configured for the targeted host 20.20.20.20, specifies the null0 interface as the route's path, and sets a tag value of 66.

```
ProCurve 9300(config)# ip route 20.20.20.20 255.255.255.255 null0 tag 66
```

On R1, a route map is configured that matches the tag 66, sets the next-hop IP address for traffic that matches tag 66 to 192.0.2.0, which is the address configured for the static null0 route on the other routers in the network. The route map also specifies a local preference value that causes the route to be the preferred route, and configures the origin as IBGP. The commands to configure this route map are as follows:

```
ProCurve 9300(config)# route-map staticb permit 1
ProCurve 9300(config-routemap staticb)# match tag 66
ProCurve 9300(config-routemap staticb)# set ip next-hop 192.0.2.0
ProCurve 9300(config-routemap staticb)# set local-preference 50
ProCurve 9300(config-routemap staticb)# set origin igp
ProCurve 9300(config-routemap staticb)# exit
```

The following commands cause static routes that match the staticb route map to be redistributed via IBGP:

```
ProCurve 9300(config)# router bgp
ProCurve 9300(config-bgp-router)# redistribute static route-map staticb
ProCurve 9300(config-bgp-router)# exit
```

Once the route update is processed, 20.20.20.20, the address of the target under attack, is installed as a null0 route on all of the routers in the network. Traffic to the targeted host is discarded.

The output of the **show ip route** command has been enhanced to display the null0 routes. For example:

```
ProCurve 9300# show ip route
Total number of IP routes: 7
B:BGP D:Connected R:RIP S:Static O:OSPF *:Candidate default
      Destination      NetMask      Gateway      Port      Cost      Type
1      3.3.3.0            255.255.255.0  0.0.0.0      v50        1          D
2      4.4.4.0            255.255.255.0  0.0.0.0      4/37       1          D
3      20.20.20.0         255.255.255.0  255.255.255.255 drop        B
4      192.100.5.0        255.255.255.0  255.255.255.255 drop        B
5      192.168.0.1        255.255.255.255 255.255.255.255 drop        1          S
6      192.168.101.0     255.255.255.0  0.0.0.0      1b1        1          D
```

IP Load Sharing for RIPv2 Routes

The IP route table can contain more than one path to a given destination. When this occurs, the device selects the path with the lowest cost as the path for forwarding traffic to the destination. If the IP route table contains more than one path to a destination and the paths each have the lowest cost, then the device uses IP load sharing to select a path to the destination. IP load sharing is also called “Equal-Cost Multi-Path (ECMP)” load sharing or just “ECMP”.

Starting in release 08.0.00, IP load sharing is supported for RIPv2 routes. In previous releases, the ProCurve device did not keep equal-cost routes from different next hops. Only the last route received for a network was kept. Starting with this release, the device stores multiple equal-cost RIPv2 routes to the same destination, and shares the traffic load among the routes.

By default, IP load sharing for RIPv2 routes is disabled. To enable it, enter the following commands:

```
ProCurve 9300(config)# router rip
ProCurve 9300(config-rip-router)# ecmp-enable
```

Syntax: [no] ecmp-enable

For more information on IP load sharing, see the *Advanced Configuration and Management Guide*.

VRRPE Slow Start Timer

In a VRRPE configuration, if a Master router goes down, the Backup router with the highest priority takes over. When the Master comes back up again, it takes over from the Backup. By default, this transition from Backup back to Master takes place immediately.

Starting in release 08.0.00, you can configure the VRRPE slow start timer feature, which causes a specified amount of time to elapse between the time the Master is restored and when it takes over from the Backup. This interval allows time for OSPF convergence when the Master is restored.

To set the VRRPE slow start timer to 30 seconds, enter the following command:

```
ProCurve 9300(config-vrrpe-router)# slow-start 30
```

Syntax: [no] slow-start <seconds>

When the VRRPE slow start timer is enabled, if the Master goes down, the Backup takes over immediately. If the Master subsequently comes back up again, the amount of time specified by the VRRPE slow start timer elapses (in this example, 30 seconds) before the Master takes over from the Backup.

The VRRPE slow start timer is effective only if another VRRPE Master (Standby) is detected. It is not effective during the initial bootup.

NOTE: The VRRPE slow start timer applies only to VRRPE configurations. It does not apply to VRRP configurations.

Dynamic Configuration of a Voice over IP (VoIP) phone

Starting in release 08.0.00, you can configure the ProCurve device to automatically detect and re-configure a VoIP phone when it is physically moved from one port to another within the same device. To do so, you must configure a **voice VLAN ID** on the port to which the VoIP phone is connected. The software stores the voice VLAN ID in the port's database for retrieval by the VoIP phone.

The dynamic configuration of a VoIP phone works in conjunction with the VoIP phone's discovery process. Upon installation, and sometimes periodically, a VoIP phone will query the ProCurve device for VoIP information, and advertise information about itself, such as, device ID, port ID, and platform. When the ProCurve device receives the VoIP phone's query, it sends the voice VLAN ID in a reply packet back to the VoIP phone. The VoIP phone then configures itself within the voice VLAN.

As long as the port to which the VoIP phone is connected has a voice VLAN ID, the phone will configure itself into that voice VLAN. If you change the voice VLAN ID, the system will immediately send the new ID to the VoIP phone, and the VoIP phone will re-configure itself with the new voice VLAN.

Configuration Notes

- This feature works with any VoIP phone that:
 - Runs CDP
 - Sends a VoIP VLAN query message
 - Can configure its voice VLAN after receiving the VoIP VLAN reply
- Automatic configuration of a VoIP phone will not work if one of the following applies:
 - You do not configure a voice VLAN ID for a port with a VoIP phone
 - You remove the configured voice VLAN ID from a port without configuring a new one
 - You remove the port from the voice VLAN
- Make sure the port is able to intercept CDP packets (**cdp run** command).
- Some VoIP phones may require a reboot after configuring or re-configuring a voice VLAN ID. For example, if your VoIP phone queries for VLAN information only once upon boot up, you must reboot the VoIP phone before it can accept the VLAN configuration. If your phone is powered by a PoE device, you can reboot the phone by disabling then re-enabling the port.

Enabling Dynamic Configuration of a Voice over IP (VoIP) phone

You can create a voice VLAN ID for a port, or for a group of ports.

To create a voice VLAN ID for a port, enter commands such as the following:

```
ProCurve 9300(config)# vlan 1001
ProCurve 9300 (config-vlan-1001)# tag e 1/2
ProCurve 9300(config-vlan-1001)# int e 1/2
ProCurve 9300(config-if-e1000-1/2)# voice-vlan 1001
```

To create a voice VLAN ID for a group of ports, enter commands such as the following:

```
ProCurve 9300(config)# vlan 1001
ProCurve 9300(config-vlan-1001)# tag e 1/1 to 1/8
ProCurve 9300(config-vlan-1001)# int e 1/1 to 1/8
ProCurve 9300(config-mif-1/1-1/8)# voice-vlan 1001
```

Syntax: [no] voice-vlan <voice-vlan-num>

where <voice-vlan-num> is a valid VLAN ID between 1 – 4095.

To remove a voice VLAN ID, use the **no** form of the command.

Viewing Voice VLAN Configurations

You can view the configuration of a voice VLAN for a particular port or for all ports.

To view the voice VLAN configuration for a port, use the **show voice-vlan <port-num>** command. The following example shows the command output results.

```
ProCurve 9300(config)# show voice-vlan ethernet 1/2
Voice vlan ID for port 1/2: 1001
```

The following example shows the message that appears when the port does not have a configured voice VLAN.

```
ProCurve 9300(config)# show voice-vlan ethernet 1/2
Voice vlan is not configured for port 1/2.
```

To view the voice VLAN for all ports, use the **show voice-vlan** command. The following example shows the command output results.

```
ProCurve 9300(config)# show voice-vlan

Port ID          Voice-vlan
1/2              1001
1/8              150
1/5              200
```

Syntax: show voice-vlan [<port-num>]

EP Layer 2 ACLs

EP Layer 2 Access Control Lists (ACLs) filter incoming traffic based on Layer 2 MAC header fields in the Ethernet/IEEE 802.3 frame. Specifically, Layer 2 ACLs filter incoming traffic based on any of the following Layer 2 fields in the MAC header:

- Source MAC address and source MAC mask
- Destination MAC address and destination MAC mask
- VLAN ID
- Ethernet type

The Layer 2 ACL feature is unique to ProCurve EP devices and differs from the existing software-based MAC address filters. MAC address filters use the CPU to filter traffic, hence, performance is limited by the CPU's processing power. Layer 2 ACLs are implemented in EP hardware and can thus filter traffic at line-rate speed.

Filtering Based on Ethertype

EP Layer 2 ACLs can filter traffic based on protocol types of a frame. Depending on the type of traffic to filter, you can select a specific Ethertype (etype) on which to filter. There are different etypes for IP and IPX traffic, which provide flexibility to filter on packet details that are beyond Layer 2. For a list of the etypes supported with Layer 2 ACLs, see "Configuring Layer 2 ACLs" on page 28.

For each Layer 2 ACL etype entry bound to a port, a Content Addressable Memory (CAM) entry is written to the corresponding CAM. You can conserve CAM space by configuring only the Layer 2 ACLs needed. For instance, to filter only IPV4-Len-5 traffic, specify that particular etype. This results in one CAM entry. Configuration examples are provided in the section "Configuring Layer 2 ACLs" on page 28.

Configuration Rules and Notes

- On ProCurve devices running **router** code, you can configure both Layer 2 ACLs and IP ACLs on the same port. Layer 2 ACLs are applied only to switched traffic, and IP ACLs are applied to routed traffic.
- You cannot bind a Layer 2 ACL to a virtual interface.
- The Layer 2 ACL feature cannot perform SNAP and LLC encapsulation type comparisons. To implement these features, use MAC address filters. You can bind MAC filters and Layer 2 ACLs on the same port, however, the device will process the traffic in software instead of in hardware.
- When MAC address filters and Layer 2 ACLs are enabled on the same port, MAC address filter processing precedes Layer 2 ACL processing; the device either forwards or drops the traffic based on the MAC filter policies, and the traffic is not subject to Layer 2 ACL processing.
- By default, when Layer 2 ACLs are enabled on a port, the device filters traffic in hardware. However, when other CPU-based features, such as Net-flow and Adaptive-Rate-Limiting are also enabled on the port, traffic is sent to the CPU for additional processing and the Layer 2 ACLs are also processed in software. Note that the performance in this case is limited by the CPU cycles.
- By default, the device processes broadcast traffic in software. Filtering of broadcast packets is not handled by the hardware.
- You can use Layer 2 ACLs to block management access to the ProCurve device. For example, you can use a Layer 2 ACL clause to block a certain host from establishing a connection to the device through Telnet.

Configuring Layer 2 ACLs

Configuring a Layer 2 ACL is similar to configuring standard and extended ACLs. Layer 2 ACL table IDs range from 400 to 499, for a maximum of 100 configurable Layer 2 ACL tables. Within each Layer 2 ACL table, you can configure from 64 (default) to 256 clauses. Each clause or entry can define a set of Layer 2 parameters for filtering. Once you completely define a Layer 2 ACL table, you must bind it to the interface for filtering to take effect.

The ProCurve device evaluates traffic coming into the port against each ACL clause. When a match occurs, the device takes the corresponding action. Once a match entry is found, the device either forwards or drops the traffic, depending upon the action specified for the clause. Once a match entry is found, the device does not evaluate the traffic against subsequent clauses.

By default, if the traffic does not match any of the clauses in the ACL table, the device drops the traffic. To override this behavior, specify a “permit any any...” clause at the end of the table to match and forward all traffic not matched by the previous clauses.

NOTE: Use precaution when placing entries within the ACL table. The Layer 2 ACL feature does not attempt to resolve conflicts and assumes you know what you are doing.

Creating a Layer 2 ACL Table

You create a Layer 2 ACL table by defining a Layer 2 ACL clause.

To create a Layer 2 ACL table, enter commands (clauses) such as the following at the Global CONFIG level of the CLI. Note that you can add additional clauses to the ACL table at any time by entering the command with the same table ID and different MAC parameters.

```
ProCurve 9300(config)# access-list 400 deny any any any etype appletalk
ProCurve 9300(config)# access-list 400 deny any any any etype ipx-raw
ProCurve 9300(config)# access-list 400 deny any any any etype ipx-snap
ProCurve 9300(config)# access-list 400 deny any any any etype ipx-llc
ProCurve 9300(config)# access-list 400 permit any any 100
```

This configuration creates a Layer 2 ACL with an ID of 400. When applied to an interface, this Layer 2 ACL table will deny all AppleTalk and IPX traffic, and permit all other traffic in VLAN 100.

For more examples of valid Layer 2 ACL clauses, see “Example Layer 2 ACL Clauses” on page 29.

Syntax: [no] access-list <num> permit | deny <src-mac> <mask> | any <dest-mac> <mask> | any [<vlan-id> | any [etype <etype-str>] [log-enable]]

The <num> parameter specifies the Layer 2 ACL table that the clause belongs to. The table ID can range from 400 to 499. You can define a total of 100 Layer 2 ACL tables.

The **permit** | **deny** argument determines the action to be taken when a match occurs.

The <src-mac> <mask> | **any** parameter specifies the source MAC address. You can enter a specific address and a comparison mask or the keyword **any** to filter on all MAC addresses. Specify the mask using F's and zeros. For example, to match on the first two bytes of the address aabb.ccdd.eeff, use the mask ffff.0000.0000. In this case, the clause matches all source MAC addresses that contain "aabb" as the first two bytes and any values in the remaining bytes of the MAC address. If you specify **any**, you don't need to specify a mask and the clause matches on all MAC addresses.

The <dest-mac> <mask> | **any** parameter specifies the destination MAC address. The syntax rules are the same as those for the <src-mac> <mask> | **any** parameter.

The optional <vlan-id> | **any** parameter specifies the vlan-id to be matched against the vlan-id of the incoming packet. You can specify **any** to ignore the vlan-id match.

The optional **etype** <etype-str> argument specifies the Ethernet type field of the incoming packet to match. The <etype-str> can be one of the following keywords:

- IPV4 (Etype=0x0800, IP version 4)
- IPV4-Len-5 (Etype=0x0800, IPV4, HeaderLen 20 bytes)
- IPV4-IGMP (Etype=0x0800, IPV4, Protocol=2)
- IPV4-IGMP-Len-5 (Etype=0x0800, IPV4-L5, Protocol=2)
- ARP (Etype=0x0806, IP ARP)
- IPX-Raw (Etype<1536, DSAP-SSAP = 0xFFFF)
- IPX-LLC (Etype<1536, DSAP-SSAP = 0xE0E0)
- IPX-SNAP (Etype<1536, DSAP-SSAP = 0xAAAA_03, Snap_Etype=0x8137)
- IPX-8137 (Etype=0x8137)
- AppleTalk (Etype<1536, DSAP-SSAP = 0xAAAA_03, Snap_Etype=0x809B)
- Apple Talk ARP (Etype<1536, DSAP-SSAP = 0xAAAA_03, Snap_Etype=0x80F3)
- Net Bios (Etype<1536, DSAP-SSAP = 0xF0F0/0xF0F1)
- IP SNAP (Etype<1536, DSAP-SSAP = 0xAAAA_03, Snap_Etype=0x0800)
- IPV6 (Etype=0x86DD, IP version 6)

The optional <log-enable> parameter enables the logging mechanism. The device accepts this command only when a **deny** clause is configured. When you enable logging for a Layer 2 ACL, all traffic matching the clause is sent to the CPU for processing and traffic is denied by the CPU. The CPU creates a log entry for the first packet that is denied and once every 10 seconds thereafter. The logging mechanism includes sending SNMP traps and log messages to the Syslog servers and writing the log entry to the log buffer on the device.

NOTE: Traffic denied by the implicit deny mechanism is not subject to logging. The implicit deny mechanism kicks in when the traffic does not match any of the clauses specified and there is no **permit any any** clause specified at the end.

Use the [no] parameter to delete the Layer 2 ACL clause from the table. When all clauses are deleted from a table, the table is automatically deleted from the system.

Example Layer 2 ACL Clauses

The following shows some examples of valid Layer 2 ACL clauses:

```
ProCurve 9300(config)# access-list 400 permit any any
```

```

ProCurve 9300(config)# access-list 400 permit any any log-enable
ProCurve 9300(config)# access-list 400 permit any any 100
ProCurve 9300(config)# access-list 400 permit any any 100 log-enable
ProCurve 9300(config)# access-list 400 permit any any
ProCurve 9300(config)# access-list 400 permit any any any log-enable
ProCurve 9300(config)# access-list 400 permit any any 100 etype ipv4
ProCurve 9300(config)# access-list 400 permit any any 100 etype ipv4 log-enable

```

Inserting and Deleting Layer 2 ACL Clauses

You can make changes to the Layer 2 ACL table definitions without unbinding and rebinding the table from an interface. For example, you can add a new clause to the ACL table, delete a clause from the table, delete the ACL table, and so on.

Binding a Layer 2 ACL Table to an Interface

To enable Layer 2 ACL filtering, bind the Layer 2 ACL table to an interface. Enter a command such as the following at the Interface level of the CLI:

```

ProCurve 9300(config)# int e 4/12
ProCurve 9300(config-int-e100-4/12)# mac access-group 400 in

```

Syntax: [no] mac access-group <num> in

The <num> parameter specifies the Layer 2 ACL table ID to bind to the interface.

Increasing the Maximum Number of Clauses per Layer 2 ACL Table

You can increase the maximum number of clauses configurable within a Layer 2 ACL table. You can specify a maximum of 256 clauses per table. The default value is 64 clauses per table.

To increase the maximum number of clauses per Layer 2 ACL table, enter a command such as the following at the Global CONFIG level of the CLI:

```

ProCurve 9300(config)# system-max l2-acl-table-entries 200

```

Syntax: system-max l2-acl-table-entries <max>

The <max> parameter specifies the maximum number of clauses per Layer 2 ACL. Enter a value from 64 to 256.

Viewing Layer 2 ACLs

Use the **show access-list** command to monitor configuration and statistics and to diagnose Layer 2 ACL tables. The following shows an example output.

```

ProCurve 9300(config)# show access-list 400

L2 MAC Access List 400:
  permit any any 100 etype ipv4
  deny any any any etype appletalk
  deny any any any etype ipx-raw
  deny any any any etype ipx-snap
  deny any any any etype ipx-llc

```

Syntax: show access-list <num>

The <num> parameter specifies the Layer 2 ACL table ID.

Layer 2 ACL-Based Rate Limiting

EP Layer 2 ACL-based rate limiting enables the ProCurve device to rate limit incoming traffic in hardware, without CPU intervention. Rate limiting in hardware enables the device to manage bandwidth at line-rate speed.

This feature is an extension to the existing IP ACL-based rate limiting on EP devices. Whereas the existing feature provides the facility to limit the rate for IP traffic that matches the permit conditions in standard or extended IP ACLs; the new feature enables you to limit traffic rates using the Layer 2 parameters defined in the associated EP Layer 2 ACL table.

In general, Layer 2 ACL-based rate limiting works along the same lines as the EP hardware-based rate limiting feature. All the rules and regulations that apply to EP rate limiting also apply to this feature. For more information about EP ACL-based rate limiting, see the *Advanced Configuration and Management Guide*.

Configuration Rules and Notes

- Layer 2 ACL-based rate limiting applies only to inbound traffic. You cannot use it to rate limit outgoing traffic.
- You can apply Layer 2 ACL-based rate limiting on a physical port. You cannot apply it to a virtual interface or a trunk port.
- You cannot use Layer 2 ACL-based filtering and Layer 2 ACL-based rate limiting on the same port. You can, however, configure one port on the device to use Layer 2 ACL-based filtering and another port on the same device to use Layer 2 ACL-based rate limiting.
- You cannot use the existing ACL-based rate limiting and Layer 2 ACL-based rate limiting on the same port. However, you can configure one port on the device to use ACL-based rate limiting and another port on the same device to use Layer 2 ACL-based rate limiting.
- On ProCurve devices running **router** code, you can use IP ACLs and Layer 2 ACL-based rate limiting on the same port. Layer 2 ACL-based rate limiting is applied only to switched traffic, and IP ACLs are applied to routed traffic.
- By default, when Layer 2 ACL-based rate limiting is enabled on the port, the device rate limits the traffic in hardware. However, when other CPU-based features, such as NetFlow and MAC filters are also enabled on the port, traffic is sent to the CPU for processing and is not subject to rate limiting.

Configuring Layer 2 ACL-Based Rate Limiting

To configure Layer 2 ACL-based rate limiting, perform the following steps:

1. Configure a Layer 2 ACL table with all the necessary clauses. See “Configuring Layer 2 ACLs” on page 28.
2. Configure a rate limit policy on a physical port using the Layer 2 ACL table ID and the desired average rate. Enter a command such as the following:

```
ProCurve 9300(config)# int e 4/25
ProCurve 9300(config-if-e1000-4/25)# rate-limit in access-group 400 10000000
```

Syntax: [no] rate-limit in access-group <acl-id> <average-rate>

The <acl-id> for Layer 2 ACLs can range from 400 to 499.

The <average-rate> is the maximum number of bits the policy allows during one second.

Editing a Layer 2 ACL Table

You can make changes to the Layer 2 ACL table definitions without unbinding and rebinding the rate limit policy. For example, you can add a new clause to the ACL table, delete a clause from the table, or delete the ACL table that is used by a rate limit policy. See “Configuring Layer 2 ACLs” on page 28.

Excluding Control Traffic from Rate Limiting

By default, the Layer 2 ACL rate limiting feature does not implement control traffic exemption for rate limiting. You can do so by configuring a respective Layer 2 ACL table and appropriate rate limit policy. For example, to prevent the ProCurve device from rate limiting OSPF control traffic, define an ACL table with the following clauses and use it to define the rate limit policy.

```
ProCurve 9300(config)# access-list 400 deny any 0000.5E00.0005 ffff.ffff.ffff any
ProCurve 9300(config)# access-list 400 deny any 0000.5E00.0006 ffff.ffff.ffff any
ProCurve 9300(config)# access-list 400 permit any any any
ProCurve 9300(config-if-e1000-4/25)# rate-limit in access-group 400 10000000
```

This configuration defines deny clauses for OSPF control packets which prevents them from being rate limited. The last clause rate limits all other traffic to 10 Mbps. Note that the deny clause will allow traffic to be forwarded without being rate limited, only if the strict ACL mode is not turned on.

Strict ACL Mode

The Layer 2 ACL rate limiting feature includes support for strict ACL mode. By default, Layer 2 ACL clauses with a **deny** action are not subject to rate limiting, and the device forwards all traffic that match these clauses in hardware. You can override this behavior by using strict ACL mode to drop the traffic that matches the deny clauses. For more about strict ACL mode, see the *Advanced Configuration and Management Guide*.

802.1s Spanning Tree Support

Multiple Spanning Tree Protocol (MSTP) as defined in IEEE 802.1s-2002 allows you to configure multiple STP instances. This ensures loop-free topology for 1 or more VLANs. Using MSTP, the entire network runs a common instance of RSTP. Within that common instance, one or more VLANs can be individually configured into distinct regions. The entire network runs the common spanning tree instance (CST) and the regions run a local instance. The local instance is known as Internal Spanning Tree (IST). The CST treats each instance of IST as a single bridge. Consequently, ports are blocked to prevent loops that might occur within an IST and also throughout the CST. In addition, MSTP can coexist with individual devices running STP or RSTP in the Common and Internal Spanning Trees instance (CIST). With the exception of the provisions for multiple instances, MSTP operates exactly like RSTP.

For example, in Figure 2 a network is configured with two regions: Region 1 and Region 2. The entire network is running an instance of CST. Each of the regions is running an instance of IST. In addition, this network contains Switch 1 running RSTP that isn't configured in a region and consequently is running in the CIST instance. In this configuration, the regions are each regarded as a single bridge to the rest of the network, as is Switch 1. The CST prevents loops from occurring across the network. Consequently, a port is blocked at either port 1/2 of region 1 switch 6, or port 3/1 of region 2 switch 4.

Additionally loops must be prevented in each of the IST instances. Within the IST Region 1, a port is blocked at port 1/2 of switch 4 to prevent a loop in that region. Within Region 2, a port is blocked at port 3/2 of switch 3 to prevent a loop in that region.

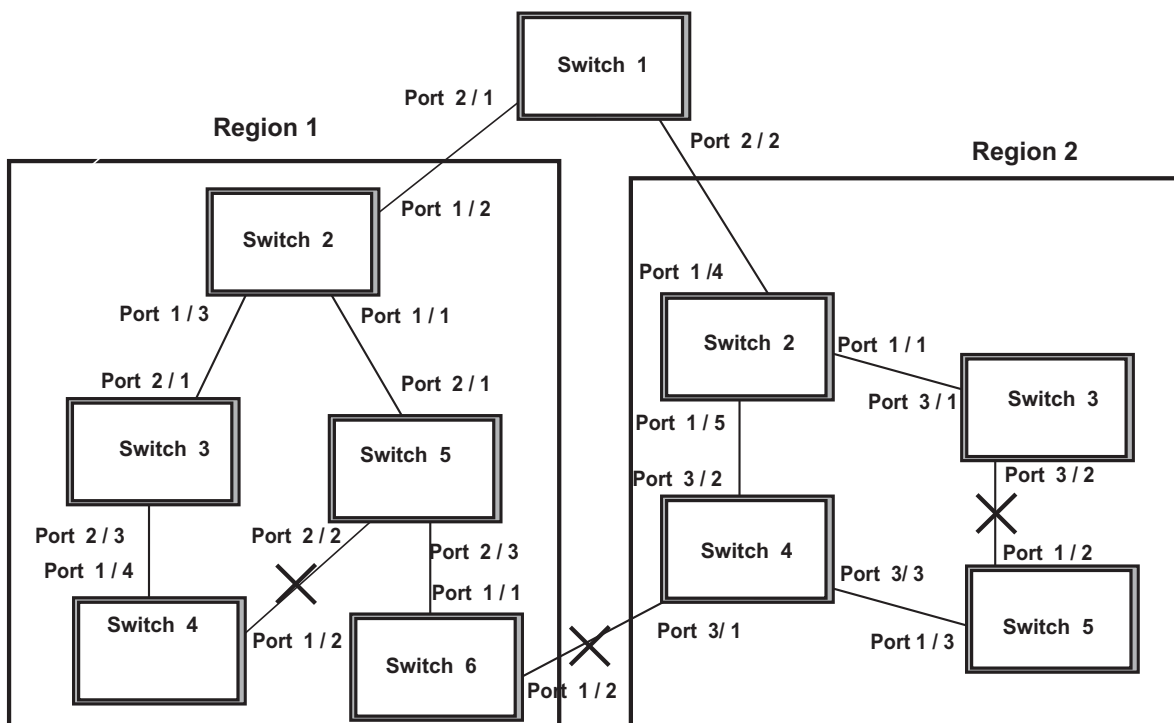


Figure 2 MSTP Configured Network

The following definitions describe the STP instances that define an MSTP configuration:

Common Spanning (CST) – MSTP runs a single instance of spanning tree across all the bridges in a network called the Common Spanning Tree (CST). This instance treats each region as a single bridge. In all other ways, it operates exactly like Rapid Spanning Tree (RSTP).

Internal Spanning Tree (IST) – The instances of spanning tree that operate within a defined region are called ISTs (Internal Spanning Tree).

Common and Internal Spanning Trees (CIST) – This is the default MSTP instance 0. It contains all of the ISTs and all bridges that are not formally configured into a region. This instance interoperates with bridges running legacy STP and RSTP implementations.

Multiple Spanning Tree Instance (MSTI) – The MSTI is identified by an MST identifier (MSTid) value between 1 and 4094. This defines an individual instance of an IST. One or more VLANs can be assigned to an MSTI. A VLAN cannot be assigned to multiple MSTIs.

Configuring MSTP

To configure a switch for MSTP, you must first configure the name and the revision on each switch that is being configured for MSTP. This name is unique to each switch. You must then create an MSTP Instance and assign an ID. VLANs are then assigned to MSTP instances. These instances must be configured on all switches that interoperate with the same VLAN assignments. Port cost, priority and global parameters can then be configured for individual ports and instances. In addition, operational edge ports and point-to-point links can be created and MSTP can be disabled on individual ports.

Each of the commands used to configure and operate MSTP are described in the following:

- “Setting the MSTP Revision Number and Name”
- “Configuring an MSTP Instance”
- “Configuring Port priority and Path cost for an MSTP Instance”
- “Configuring Priority for an MSTP Instance”
- “Setting the MSTP Global Parameters”
- “Setting Ports To Be Operational Edge Ports”
- “Setting Point-to-Point Link”
- “Committing MSTP Configuration ID Changes”
- “Disabling MSTP on a Port”
- “Forcing Ports to Transmit the Configured BPDU Version”
- “Enabling MSTP on a Switch”

Setting the MSTP Revision Number and Name

Each switch that is running MSTP is configured with a name and revision number. These apply to the switch which can have many different VLANs that can belong to many different MSTP regions.

To configure an MSTP name and revision number, use a command such as the following at the Global Configuration level:

```
ProCurve 9300(config)# mstp configuration name procurve revision 4
```

Syntax: [no] mstp configuration name <name> revision <revision-number>

The **name** parameter defines an ASCII name for the MSTP configuration. The default name is the MAC address of the switch expressed as a string.

The **revision** parameter specifies the revision level for MSTP that you are configuring on the switch. It can be a number between 0 and 65535.

Configuring an MSTP Instance

An MSTP instance is configured with an MSTP ID for each region. Each region can contain one or more VLANs. To configure an MSTP instance and assign a range of VLANs, use a command such as the following at the Global Configuration level:

```
ProCurve 9300(config)# mstp configuration instance 7 vlan 4 to 7
```

Syntax: [no] mstp configuration instance <instance-number> vlan <number>

The **instance** parameter defines the number for the instance of MSTP that you are configuring.

The **vlan** parameter assigns one or more VLANs or a range of VLANs to the instance defined in this command.

Configuring Port priority and Path cost for an MSTP Instance

Port priority and path cost can be configured for a specified instance. To configure an MSTP instance, use a command such as the following at the Global Configuration level:

```
ProCurve 9300(config)# mstp instance 1 ethernet 3/1 priority 32 path-cost 200
```

Syntax: [no] mstp instance <instance-number> [ethernet <slot/port> priority <port-priority>] path-cost <cost>

The <instance-number> variable is the number of the instance of MSTP that you are configuring port priority and path cost for.

The **ethernet** <slot/port> parameter configures a port's operation in this instance.

You can set a **priority** to the instance that gives it forwarding preference over lower priority instances within a VLAN or on the switch. A higher number for the priority variable means a lower forwarding priority. Acceptable values are 0 - 240 in increments of 16. The default value is 128.

A **path-cost** can be assigned to a port to bias traffic towards or away from a path during periods of rerouting. Possible values are 1 - 200000000. The default values for Ethernet interfaces are: 2,000,000 for 10 Mbps, 200,000 for 100 Mbps, 20,000 for Gigabit Ethernet, and 2,000 for 10 Gigabit Ethernet.

Configuring Priority for an MSTP Instance

Priority can be configured for a specified instance. To configure priority for an MSTP instance, use a command such as the following at the Global Configuration level:

```
ProCurve 9300(config)# mstp instance 1 priority 32768
```

Syntax: [no] mstp instance <instance-number> priority <port-priority>

The <instance-number> variable is the number for the instance of MSTP that you are configuring.

You can set a **priority** to the instance that gives it forwarding preference over lower priority instances within a VLAN or on the switch. A higher number for the priority variable means a lower forwarding priority. Acceptable values are 0 - 61440 in increments of 4096. The default value is 32768.

Setting the MSTP Global Parameters

MSTP has many of the options available in RSTP as well as some unique options. The following command is used to configure MSTP Global parameters for all instances configured on a switch:

```
ProCurve 9300(config)# mstp force-version 0 forward-delay 10 hello-time 4 max-age 8 max-hops 9
```

Syntax: [no] mstp force-version <mode-number> forward-delay <value> hello-time <value> max-age <value> max-hops <value>

The **force-version** parameter forces the bridge to send BPDUs in a specific format. You can specify one of the following <mode-number> values:

- 0 – The STP compatibility mode. Only STP BPDUs will be sent.
- 2 – The RSTP compatibility mode. Only RSTP BPDUS will be sent.
- 3 – MSTP mode. In this default mode, only MSTP BPDUS will be sent.

The **forward-delay** <value> specifies how long a port waits before it forwards an RST BPDUS after a topology change. This can be a value from 4 – 30 seconds. The default is 15 seconds.

The **hello-time** <value> parameter specifies the interval between two hello packets. The parameter can have a value between 1 – 10 seconds. The default is 2 seconds.

The **max-age** <value> parameter specifies the amount of time the device waits to receive a hello packet before it initiates a topology change. You can specify a value between 6 – 40 seconds. The default value is 20 seconds.

The **max-hops** <value> parameter specifies the maximum hop count. You can specify a value between 1 – 40 hops. The default value is 20 hops.

Setting Ports To Be Operational Edge Ports

You can define specific ports as edge ports for the region that they are configured in to connect to devices that are not running STP, RSTP or MSTP such as a host. If a port is connected to an end device such as a PC, the port can be configured as an edge port. The following command is used to configure ports as operational edge ports:

```
ProCurve 9300(config)# mstp admin-edge-port ethernet 3/1
```

Syntax: [no] mstp admin-edge-port ethernet <slot/port>

The <slot/port> parameter specifies a port or range of ports as edge ports in the instance they are configured in.

Setting Point-to-Point Link

You can set a point-to-point link between ports to increase the speed of convergence. To create a point-to-point link between ports, use a command such as the following at the Global Configuration level:

```
ProCurve 9300(config)# mstp admin-pt2pt-mac ethernet 2/5 ethernet 4/5
```

Syntax: [no] mstp admin-pt2pt-mac ethernet <slot/port>

The <slot/port> parameter specifies a port or range of ports to be configured for point-to-point links to increase the speed of convergence.

Committing MSTP Configuration ID Changes

You need to use the **mstp commit** command only if changes are made to an MSTP configuration identifier while MSTP is operational. If MSTP is not operational, any configuration ID changes are applied immediately and there's no need to use "commit" command.

After you make MSTP configuration identifier changes, you must use the following command at the Global Configuration level to commit the change:

```
ProCurve 9300(config)# mstp commit
```

Syntax: [no] mstp commit

Disabling MSTP on a Port

To disable MSTP on a specific port, use a command such as the following at the Global Configuration level:

```
ProCurve 9300(config)# mstp disable 2/1
```

Syntax: [no] mstp disable <slot/port>

The <slot/port> variable is location of the port that you want to disable MSTP for.

Forcing Ports to Transmit the Configured BPDU Version

To force a port to transmit a BPDU of the version that is configured on this device (3 is default), use a command such as the following at the Global Configuration level:

```
ProCurve 9300(config)# mstp force-migration-check ethernet 3/1
```

Syntax: [no] mstp force-migration-check ethernet <slot/port>

The <slot/port> variable specifies the port or ports that you want to transmit an MSTP BPDU from.

Enabling MSTP on a Switch

To enable MSTP on your switch, use a command such as the following at the Global Configuration level:

```
ProCurve 9300(config)# mstp start
```

Syntax: [no] mstp start

Sample Configuration

In the example shown in “Sample MSTP Configuration, below, four ProCurve switch routers are configured in two regions. There are four VLANs in four instances in Region 2 and Region 1 is in the CIST.




Figure 3 SAMPLE MSTP Configuration

RTR1 Configuration

```
ProCurve 9300(config)# mstp config name Reg1 revision 1
ProCurve 9300(config)# mstp admin-pt2pt-mac ethe 10/1 to 10/2
ProCurve 9300(config)# mstp start
ProCurve 9300(config)# hostname RTR1
```

Core 1 Configuration

```
ProCurve 9300(config)# trunk switch ethernet 2/9 to 2/12
ProCurve 9300(config)# trunk switch ethernet 2/13 to 2/14
ProCurve 9300(config)# vlan 1 name DEFAULT-VLAN by port
ProCurve 9300(config-vlan-1)# no spanning-tree
ProCurve 9300(config-vlan-1)# exit
ProCurve 9300(config)# vlan 20 by port
ProCurve 9300(config-vlan-20)# tagged ethernet 2/9 to 2/14 ethernet 2/16
ProCurve 9300(config-vlan-20)# no spanning-tree
ProCurve 9300(config-vlan-20)# exit
ProCurve 9300(config)# vlan 21 by port
ProCurve 9300(config-vlan-21)# tagged ethernet 2/9 to 2/14 ethernet 2/16
ProCurve 9300(config-vlan-21)# no spanning-tree
ProCurve 9300(config-vlan-21)# exit
ProCurve 9300(config)# vlan 22 by port
ProCurve 9300(config-vlan-22)# tagged ethernet 2/9 to 2/14 ethernet 2/16
ProCurve 9300(config-vlan-22)# no spanning-tree
ProCurve 9300(config-vlan-22)# exit
ProCurve 9300(config)# mstp config name HR revision 2
ProCurve 9300(config)# mstp config instance 20 vlan 20
ProCurve 9300(config)# mstp config instance 21 vlan 21
ProCurve 9300(config)# mstp config instance 22 vlan 22
ProCurve 9300(config)# mstp instance 0 priority 8192
ProCurve 9300(config)# mstp admin-pt2pt-mac ethernet 2/9 to 2/14
ProCurve 9300(config)# mstp admin-pt2pt-mac ethernet 2/16
ProCurve 9300(config)# mstp start
ProCurve 9300(config)# hostname CORE1
```

Core2 Configuration

```
ProCurve 9300(config)# trunk switch ethernet 3/5 to 3/6
ProCurve 9300(config)# trunk switch ethernet 3/17 to 3/20
ProCurve 9300(config)# vlan 1 name DEFAULT-VLAN by port
ProCurve 9300(config-vlan-1)# no spanning-tree
ProCurve 9300(config-vlan-1)# exit
ProCurve 9300(config)# vlan 20 by port
ProCurve 9300(config-vlan-20)# tagged ethernet 3/5 to 3/6 ethernet 3/17 to 3/20
ProCurve 9300(config-vlan-20)# no spanning-tree
ProCurve 9300(config-vlan-20)# exit
ProCurve 9300(config)# vlan 21 by port
ProCurve 9300(config-vlan-21)# tagged ethernet 3/5 to 3/6 ethernet 3/17 to 3/20
ProCurve 9300(config-vlan-21)# no spanning-tree
ProCurve 9300(config-vlan-21)# exit
ProCurve 9300(config)# vlan 22 by port
ProCurve 9300(config-vlan-22)# tagged ethe 3/5 to 3/6 ethe 3/17 to 3/20
ProCurve 9300(config-vlan-22)# no spanning-tree
ProCurve 9300(config-vlan-22)# exit
ProCurve 9300(config)# mstp config name HR revision 2
ProCurve 9300(config)# mstp config instance 20 vlan 20
ProCurve 9300(config)# mstp config instance 21 vlan 21
ProCurve 9300(config)# mstp config instance 22 vlan 22
ProCurve 9300(config)# mstp admin-pt2pt-mac ethernet 3/17 to 3/20 ethernet 3/5
to 3/6
ProCurve 9300(config)# mstp admin-pt2pt-mac ethernet 3/10
ProCurve 9300(config)# mstp start
ProCurve 9300(config)# hostname CORE2
```

LAN 4 Configuration

```
ProCurve 9300(config)# trunk switch ethernet 3/5 to 3/6
ProCurve 9300(config)# trunk switch ethernet 3/1 to 3/2
ProCurve 9300(config)# vlan 1 name DEFAULT-VLAN by port
ProCurve 9300(config-vlan-1)# no spanning-tree
ProCurve 9300(config-vlan-1)# exit
ProCurve 9300(config)# vlan 20 by port
ProCurve 9300(config-vlan-20)# tagged ethernet 3/1 to 3/2 ethernet 3/5 to 3/6
ProCurve 9300(config-vlan-20)# no spanning-tree
ProCurve 9300(config-vlan-20)# exit
ProCurve 9300(config)# vlan 21 by port
ProCurve 9300(config-vlan-21)# tagged ethernet 3/1 to 3/2 ethe 3/5 to 3/6
ProCurve 9300(config-vlan-21)# no spanning-tree
ProCurve 9300(config-vlan-21)# exit
ProCurve 9300(config)# vlan 22 by port
ProCurve 9300(config-vlan-22)# tagged ethernet 3/1 to 3/2 ethe 3/5 to 3/6
ProCurve 9300(config-vlan-22)# no spanning-tree
ProCurve 9300(config-vlan-22)# exit
ProCurve 9300(config)# mstp config name HR revision 2
ProCurve 9300(config)# mstp config instance 20 vlan 20
ProCurve 9300(config)# mstp config instance 21 vlan 21
ProCurve 9300(config)# mstp config instance 22 vlan 22
ProCurve 9300(config)# mstp admin-pt2pt-mac ethernet 3/5 to 3/6 ethernet 3/1 to 3/2
ProCurve 9300(config)# mstp start
ProCurve 9300(config)# hostname LAN4
```

Displaying MSTP Information

MSTP information can be displayed using the commands shown below.

```
ProCurve 9300# show mstp

Multiple Spanning Tree Protocol (IEEE 802.1s) is OPERATIONAL
Configuration name: [HR], revision level: 2
Total 4 MST instance(s):
  CIST 0 mapped vlans: 1 to 19 36 to 4094
  MSTI 20 mapped vlans: 20
  MSTI 21 mapped vlans: 21
  MSTI 22 mapped vlans: 22
  Disabled ports:
Force version: 3
Configured hello time 2 forward delay 15 max age 20 max hops 20
Operational hello time 2 forward delay 15 max age 20 max hops 19
```

Syntax: show mstp [<instance-number> | ethernet <slot/port> | pending-changes

The <instance-number> variable specifies the MSTP instance that you want to display information for.

The **ethernet** parameter is used to specify a port that you want to display MSTP information for.

The **pending-changes** parameter specifies that you want to display pending MSTP configuration identifier changes.

Table 8: Output from Show MSTP

This Field...	Displays...
Configuration Name	The MSTP configuration name that is configured on the switch.
Revision Level	The MSTP revision level that is configured on the switch.
Total MST Instances	These fields display the number of MST instances operating on the switch, and the VLANs that reside in the CST and each MST instance.
Disabled Ports	The ethernet ports that have been disabled from MSTP.
Force Version	The format that BPDUs are sent in:
Configured:	
Hello Time	The hello value configured on this device.
Forward Delay	The period of time that the device will wait before it forwards a an RST BPDU as configured on this device.
Max Age	The interval that this device will wait for a hello packet before initiating a topology change as configured.
Max Hops	The number of hops between bridges before a packet times out and is stops circulating as configured on this device.
Operational:	
Hello Time	The hello value as received from the root bridge. If this is the root bridge, the value will be equal to the Configured Hello Time.

Table 8: Output from Show MSTP (Continued)

This Field...	Displays...
Forward Delay	The forward delay value as received from the root bridge. If this is the root bridge, the value will be equal to the Configured Forward Delay.
Max Age	The max age value as received from the root bridge. If this is the root bridge, the value will be equal to the Configured Max Age.
Max Hops	The max hops value as received from the root bridge. If this is the root bridge, the value will be equal to the Configured Max Hops.

Displaying MSTP Information for a Specified Instance

The following example displays MSTP information specified for an MSTP instance.

```
ProCurve 9300# show mstp 22

##### MST 22 vlans mapped: 22
Member ports:   ethe 2/9 to 2/14 ethe 3/1 to 3/2
Bridge ID:      address 00e0.52aa.2f00 priority 32790 (32768 sysid 22)
Reg Root ID:    address 00e0.52c2.cf60 priority 22 (0 sysid 22)
                  port    2/9      int path cost: 200000
Remaining hops: 19

Interface Role  State Designated ID      Ext-cost Int-cost Prio.Num Pt2Pt Edge
2/9       Root  FWD  001600e052c2cf60 --      200000 128.73  Y    N
2/10      Altr  BLK  001600e052c2cf60 --      200000 128.74  Y    N
2/11      Altr  BLK  001600e052c2cf60 --      200000 128.75  Y    N
2/12      Altr  BLK  001600e052c2cf60 --      200000 128.76  Y    N
2/13      Altr  BLK  8016000cdb310400 --      200000 128.77  Y    N
2/14      Altr  BLK  8016000cdb310400 --      200000 128.78  Y    N
2/16      Altr  BLK  801600e05202dc00 --      20000  128.129 Y    N
```

Syntax: show mstp <number>

Table 9: Output from Show MSTP Instance

This Field...	Displays...
Member Ports	Ports that are included in the specified MSTP instance.
Bridge ID	The ID of the bridge.
Reg Root ID	The Reg Root ID is the MAC address of the Root Bridge for the local region. This value is "self" if the device you are accessing is the Root Bridge for the local region.
Remaining Hops	Hops remaining valid for this message in relationship to the regional root bridge.
Interface	The port number of the interface.

Table 9: Output from Show MSTP Instance (Continued)

This Field...	Displays...
Role	The current role of the port: <ul style="list-style-type: none"> • Master • Root • Designated • Alternate • Backup
State	The port's current 802.1w state. A port can have one of the following states: <ul style="list-style-type: none"> • Forwarding • Learning • Blocking
Designated ID	The ID of the bridge that sent the best BPDU that was received on this port.
Int-cost	The configured path cost on a link connected to this port within the internal MSTP region.
Prio.Num	The configured priority and number of the port. The default value for the priority is 128.
Pt2Pt	Indicates if the port is configured with a point-to-point link: <ul style="list-style-type: none"> • Y – The port is configured in a point-to-point link • N – The port is not configured in a point-to-point link
Edge	Indicates if the port is configured as an operational edge port: <ul style="list-style-type: none"> • Y – indicates that the port is defined as an edge port. • N – indicates that the port is not defined as an edge port

Displaying MSTP Information for CIST Instance 0

Instance 0 is the Common and Internal Spanning Tree Instance (CIST). When you display information for this instance there are some differences with displaying other instances. The following example displays MSTP information for CIST Instance 0.

```
ProCurve# show mstp 0

##### MST 0 vlans mapped: 1 to 4094
Member ports:  ethe 1/1 to 1/18 ethe 1/20 to 1/28 ethe 3/1 to 3/16 ethe 4/1 to 4/2
               ethe 5/1 to 5/48 ethe 7/1 to 7/48
Bridge ID:      address 0004.8089.3300 priority 0 (0 sysid 0)
Root ID:        self
Reg Root ID:    self
Message age: 0  remaining hops: 20

Interface Role  State Designated ID      Ext-cost Int-cost Prio.Num Pt2Pt Edge
1/10      Desg  FWD  0000000480893300  200000  200000  128.10  Y   N
1/11      Desg  FWD  0000000480893300  200000  200000  128.11  Y   N
1/12      Desg  FWD  0000000480893300  200000  200000  128.12  Y   N
1/13      Desg  FWD  0000000480893300  200000  200000  128.13  N   N
4/1       Desg  FWD  0000000480893300  2000    2000    128.193 Y   N
```

Table 10 contains the display parameters that differ from those used for instances other than instance 0. All other parameters are described in Table 9.

Table 10: Output from Show MSTP Instance for the 0 Instance

This Field...	Displays...
Root ID	In the 0 instance, the Root ID is the MAC address of the Root Bridge for the CIST. This value is "self" if the device you are accessing is the Root Bridge for the CIST.
Reg Root ID	In the 0 instance, the Reg Root ID is the MAC address of the Root Bridge for the local region. This value is "self" if the device you are accessing is the Root Bridge for the local region.
Message age	The amount of the max age value that the message used in transit from the root bridge external to the local region.
Ext-Cost	The configured path cost on a link connected to this port to an external MSTP region.

Creating an Alias for a CLI Command

Starting in release 08.0.00, you can create *aliases* for CLI commands. An alias serves as a shorthand version of a longer CLI command. For example, you can create an alias called **shoro** for the CLI command **show ip route**. Then when you enter **shoro** at the command prompt, the **show ip route** command is executed.

To create an alias called **shoro** for the CLI command **show ip route**, enter the following command:

```
ProCurve 9300(config)# alias shoro = show ip route
```

Syntax: [no] alias <alias-name> = <cli-command>

The <alias-name> must be a single word, without spaces.

After the alias is configured, entering **shoro** at either the Privileged EXEC or CONFIG levels of the CLI, executes the **show ip route** command.

To create an alias called **wrsbc** for the CLI command **copy running-config tftp 10.10.10.10 test.cfg**, enter the following command:

```
ProCurve 9300(config)# alias wrsbc = copy running-config tftp 10.10.10.10 test.cfg
```

To remove the **wrsbc** alias from the ProCurve device's configuration, enter one of the following commands:

```
ProCurve 9300(config)# no alias wrsbc = copy running-config tftp 10.10.10.10 test.cfg
```

or

```
ProCurve 9300(config)# unalias wrsbc
```

Syntax: unalias <alias-name>

The specified <alias-name> must be the name of an alias already configured on the ProCurve device.

To display the aliases currently configured on the ProCurve device, enter the following command at either the Privileged EXEC or CONFIG levels of the CLI:

```
ProCurve 9300# alias
      wrsbc    copy running-config tftp 10.10.10.10 test.cfg
      shoro    show ip route
```

Syntax: alias

Notes

- You cannot include additional parameters with the alias at the command prompt. For example, after you create the **shoro** alias, **shoro bgp** would not be a valid command.
- If configured on the ProCurve device, authentication, authorization, and accounting is performed on the actual command, not on the alias for the command.
- To save an alias definition to the device's startup-config file, use the **write memory** command.

Directing Debugging Output to Multiple Destinations

By default, debugging output (output generated by **debug** commands) is directed only to the console. You can optionally direct debugging output to other destinations, including the Syslog buffer, or a specified Telnet or SSH session.

In previous releases, debugging output could be directed only to a single destination. Starting in release 08.0.00, you can direct debugging output to multiple destinations. This allows debugging output to be displayed on multiple user sessions concurrently, which can be useful when multiple engineers are troubleshooting a problem from multiple sites.

You can send debugging output to all destinations, or to specified destinations. In addition, you can discontinue sending debugging output to specified destinations, without affecting the debugging output sent to other destinations. In previous releases, if multiple users were using **debug** commands, changing the destination for debugging output on one user's session changed the destination for debugging output for all user sessions.

To send debugging output to the console, the Syslog buffer, and all currently active Telnet and SSH sessions on the ProCurve device, enter the following command:

```
ProCurve 9300# debug destination all
```

To stop sending debugging output to all destinations, enter the following command:

```
ProCurve 9300# no debug destination all
```

Syntax: [no] debug destination console | logging | telnet <num> | ssh <num> | all

When debugging output is being directed to all destinations, if you then want to stop sending debugging output to Telnet session 1, but keep sending debugging output to all of the other destinations, enter the following command:

```
ProCurve 9300# no debug destination telnet 1
```

NOTE: To determine the number of your Telnet or SSH session, use the **show who** command.

Similarly, if you want to stop sending debugging output to the Syslog buffer but not to the other destinations, enter the following command:

```
ProCurve 9300# no debug destination logging
```

sFlow Enhancements

sFlow is a system for observing traffic flow patterns and quantities within and among a set of network devices. When sFlow is configured and enabled on a ProCurve device, the device samples packet flows, gathers information about the sampled traffic, and exports the information to external devices known as sFlow collectors.

sFlow is described in RFC 3176, "InMon Corporation's sFlow: A Method for Monitoring Traffic in Switched and Routed Networks". For information on configuring sFlow on ProCurve devices, see the "sFlow" section in the *Advanced Configuration and Management Guide*.

Release 08.0.00 includes the following enhancements to ProCurve's support for sFlow:

- Support for sFlow version 5
- Support for selecting which and how often data destined to the CPU (for example, Telnet sessions) is sent to the sFlow collector

Support for sFlow Version 5

sFlow version 5 modifies and enhances the format of the data sent to the sFlow collector. The new version defines a new datagram syntax for the sFlow agent to report flow samples and interface counters to sFlow collectors, as well as a number of new features.

In previous releases, ProCurve devices exported data in sFlow version 2 format. Starting with release 08.0.00, ProCurve devices can export data in either sFlow version 5 or version 2 format. sFlow version 5 collectors are compatible with both sFlow version 2 and version 5 agents running inside a ProCurve device.

Starting in release 08.0.00, the sFlow agent exports sFlow version 2 flow samples by default. You can optionally configure the sFlow agent to export sFlow version 5 flow samples.

Release 08.0.00 includes the following features related to sFlow version 5:

- Support for the sFlow version 5 datagram
- Support for the sFlow version 5 MIB
- Support for sub-agents
- Configurable sFlow export packet size
- Support for the new data field and sample type length in flow samples
- Configurable interval for exporting ProCurve-specific data structure

New commands have been added to the CLI to allow you to do the following:

- Specify the sFlow version used for exporting sFlow data
- Specify the maximum flow sample size
- Export CPU and memory usage information to the sFlow collector
- Specify the polling interval for exporting CPU and memory usage information to the sFlow collector

Specifying the Version Used for Exporting sFlow Data

Starting in release 08.0.00, by default the sFlow agent on the ProCurve device exports sFlow data in version 2 format. You can optionally change this setting so that the sFlow agent on the ProCurve device exports data in version 5 format.

To do this, enter the following command:

```
ProCurve 9300(config)# sflow version 5
```

Syntax: [no] sflow version 2 | 5

The default is 2.

Specifying the Maximum Flow Sample Size

You can specify the maximum size in bytes of the flow sample sent to the sFlow collector. If a packet is larger than the specified maximum size, then only the contents of the packet up to the specified maximum number of bytes is exported. If the size of the packet is smaller than the specified maximum, then the entire packet is exported.

To specify 1000 bytes as the maximum flow sample size, enter the following command:

```
ProCurve 9300(config)# sflow max-packet-size 1000
```

Syntax: [no] sflow max-packet-size <size>

For both sFlow version 2 and version 5, the default maximum flow sample size is 256 bytes.

For sFlow version 5, the maximum flow sample size is 1000 bytes.

Exporting CPU and Memory Usage Information to the sFlow Collector

Starting in release 08.0.00, you can optionally configure the sFlow agent on the ProCurve device to export information about CPU and memory usage to the sFlow collector.

To export CPU usage and memory usage information, enter the following command:

```
ProCurve 9300(config)# sflow export system-info
```

Syntax: [no] sflow export system-info

By default, CPU usage information and memory usage information are not exported.

Specifying the Polling Interval for Exporting CPU and Memory Usage Information to the sFlow Collector

The polling interval defines how often sFlow data for a port is sent to the sFlow collector. You can optionally set the polling interval used for exporting CPU and memory usage information.

For example, to set the polling interval for exporting CPU and memory usage information to 30 seconds, enter the following command:

```
ProCurve 9300(config)# sflow export system-info 30
```

Syntax: [no] sflow export system-info <seconds>

You can specify a polling interval from 5 seconds to 1,800 seconds (30 minutes). The default polling interval for exporting CPU and memory usage information is 300 seconds (5 minutes).

Exporting CPU-Directed Data to the sFlow Collector

Starting in release 08.0.00, you can select which and how often data destined to the CPU (for example, Telnet sessions) is sent to the sFlow collector.

New commands have been added to the CLI to allow you to do the following:

- Enable the sFlow agent to export CPU-directed data
- Specify the sampling rate for exported CPU-directed data

Enabling the sFlow Agent to Export CPU-Directed Data

To enable the sFlow agent on a ProCurve device to export data destined to the CPU to the sFlow collector, enter the following command:

```
ProCurve 9300(config)# sflow export cpu-traffic
```

Syntax: [no] sflow export cpu-traffic

By default, this command is disabled. The sFlow agent does not send data destined to the CPU to the sFlow collector.

Specifying the Sampling Rate for Exported CPU-Directed Data

The sampling rate is the average ratio of the number of packets incoming on an sFlow-enabled port, to the number of flow samples taken from those packets. You can optionally set the sampling rate for CPU-directed data exported to the sFlow collector. For example, to set this sampling rate to 2048, enter the following command:

```
ProCurve 9300(config)# sflow export cpu-traffic 2048
```

Syntax: [no] sflow export cpu-traffic <rate>

The default sampling rate depends on the ProCurve device being configured. See “Changing the Sampling Rate” in the “sFlow” section of the *Advanced Configuration and Management Guide* for the default sampling rate for each kind of ProCurve device.

Jumbo Packet Counter in show Command Output

Starting in release 08.0.00, the output of the **show interfaces ethernet** and **show statistics ethernet** command has been enhanced to display information about the number of packets received that were longer than 1518 octets. In the following examples, the new output is highlighted in bold.

```
ProCurve 9300# show interfaces ethernet e 1/1
FastEthernet1/1 is up, line protocol is up
  Hardware is FastEthernet, address is 0004.8085.c500 (bia 0004.8085.c500)
  Configured speed auto, actual 100Mbit, configured duplex fdx, actual fdx
  Member of L2 VLAN ID 1, port is untagged, port state is FORWARDING
  STP configured to ON, priority is level0, flow control enabled
  mirror disabled, monitor disabled
  Not member of any active trunks
  Not member of any configured trunks
  No port name
  Internet address is 10.1.1.4/24, MTU 1518 bytes, encapsulation ethernet
  300 second input rate: 0 bits/sec, 0 packets/sec, 0.00% utilization
  300 second output rate: 0 bits/sec, 0 packets/sec, 0.00% utilization
  0 packets input, 0 bytes, 0 no buffer
  Received 0 broadcasts, 0 multicasts, 0 unicasts
  0 input errors, 0 CRC, 0 frame, 0 ignored
  0 filtered, 0 runts, 0 giants, DMA received 0 packets, 0 jumbos
  4 packets output, 256 bytes, 0 underruns
  Transmitted 4 broadcasts, 0 multicasts, 0 unicasts
  0 output errors, 0 collisions, DMA transmitted 4 packets, 0 jumbos
```

Syntax: show interface ethernet [<slot>/]<portnumber>

```

ProCurve 9300# show statistics ethernet e 1/25
Port Link State Dupl Speed Trunk Tag Priori MAC Name
1/1 Up Forward Full 100M None No level0 0004.8085.c500

Port 1/1 Counters:
      InOctets          0          OutOctets          256
      InPkts           0          OutPkts            4
      DMA recvd        0          DMA xmitted        4
      InBroadcastPkts 0          OutBroadcastPkts   4
      InMulticastPkts 0          OutMulticastPkts   0
      InUnicastPkts   0          OutUnicastPkts     0
      InJumboPkts      0          OutJumboPkts      0
      InDiscards       0          OutDiscards        0
      InErrors         0          OutErrors          0
      InCollisions     0          OutCollisions      0
                                   OutLateCollisions    0
      Alignment        0          FCS                0
      GiantPkts        0          ShortPkts          0
      InBitsPerSec     0          OutBitsPerSec      0
      InPktsPerSec     0          OutPktsPerSec      0
      InUtilization    0.00%    OutUtilization     0.00

```

Syntax: show statistics ethernet

Displaying VLANs in Numerical Order

Starting in release 08.0.00, the output of the **show run** command displays the VLANs configured on the ProCurve device in numerical order (for example, VLAN 100, VLAN 200, and so on). In previous releases, the **show run** command displayed the VLANs in the order they were configured on the device.

Specifying a Host Name in an ACL Statement

Starting in release 08.0.00, you can specify a host name within the configuration of an access control list. Previous releases supported only IP addresses in ACL statements.

For example, the following are valid ACL statements in release 08.0.00:

```

access-list 101 deny ip host www.google.com host www.ibm.com
access-list 102 permit tcp host www.sbc.com any eq telnet log

```

This enhancement is valid for both rule-based and flow-based ACLs.

When you enter an ACL statement that contains a host name in the CLI, the ProCurve device attempts to resolve the host name using the DNS resolver. The ProCurve device checks its DNS cache for an entry corresponding to the host name; if an entry is not found, the device sends a DNS query to the configured DNS server. When the host name is resolved to an IP address, a Layer 4 CAM entry is created for the ACL.

If the host name cannot be resolved, then the ACL statement is not activated. When you enter the **show run** command, a line such as the following appears for ACL statements referring to hosts that aren't resolved:

```

permit udp host 3.3.3.3 host www.yahoo.com (not in effect)

```

When a host name in an ACL statement cannot be resolved, the ProCurve device will periodically attempt to resolve it. In addition, the ProCurve device will periodically attempt to resolve host names that had been resolved previously, but are no longer resolved because TTL expired. If the resolution is successful and the IP address has changed, then the ACL's Layer 4 CAM entry is updated; otherwise, it remains unchanged.

Notes

- In order for the host name ACL feature to work, DNS must be properly configured on the ProCurve device, and the configured DNS server must be reachable from the ProCurve device.

- If the host name times out, the ProCurve device attempts to resolve the host name again. If the resolution is successful, there are two possibilities:
 - When the host-name-to-IP-address mapping is unchanged, the Layer 4 CAM entry remains intact.
 - When the host-name-to-IP-address mapping changes, the Layer 4 CAM entry is flushed with the new IP address for the host.

If the host name resolution is unsuccessful, the Layer 4 CAM entry for the ACL is removed. The output of the **show run** command displays the ACL entry as "not in effect". For example, if the host name `www.procurve.com` is not resolved, then the output of the **show run** command would have the following line:

```
permit udp host 3.3.3.3 host www.procurve.com (not in effect)
```

The ProCurve device will then make multiple attempts to resolve the host name. If and when the host name is resolved, a new Layer 4 CAM entry is created for the newly resolved IP address.

- When the ProCurve device is booted, it may take a few seconds for an ACL statement containing a host name to take effect. This is because the device must first resolve the host name and create a Layer 4 CAM entry. During this brief period, the output of the **show run** command displays the ACL entry as "not in effect".

Configuring a Domain Name List and Using Domain Look Up

The Domain Name Server (DNS) resolver is a feature that sends and receives queries to and from the DNS server on behalf of a client. Prior to this release, the feature lets you use one domain name to perform Telnet, ping, traceroute and other DNS query commands. You define one domain name on a ProCurve routing switch and up to four DNS servers. Host names and their IP addresses are configured on the DNS servers.

When a client performs a DNS query, all hosts within that domain can be recognized. After you define a domain name, the ProCurve routing switch automatically appends the appropriate domain to a host and forwards it to the DNS servers for resolution.

For example, if the domain "eng.company.com" is defined on a ProCurve routing switch and you want to initiate a ping to "mary". You need to reference only the host name instead of the host name and its domain name. For example, you could enter the following command to initiate the ping:

```
ProCurve 9300>ping mary
```

The ProCurve routing switch qualifies the host name by appending a domain name, for example, `mary.eng.company.com`. This qualified name is sent to the DNS server for resolution. If there are four DNS servers configured, it is sent to the first DNS server. If the host name is not resolved, it is sent to the second, and

then the third, then the fourth DNS server. If a match is found, a response is sent back to the client with the host's IP address. If no match is found, a "unknown host" message is returned. (See Figure 4.)

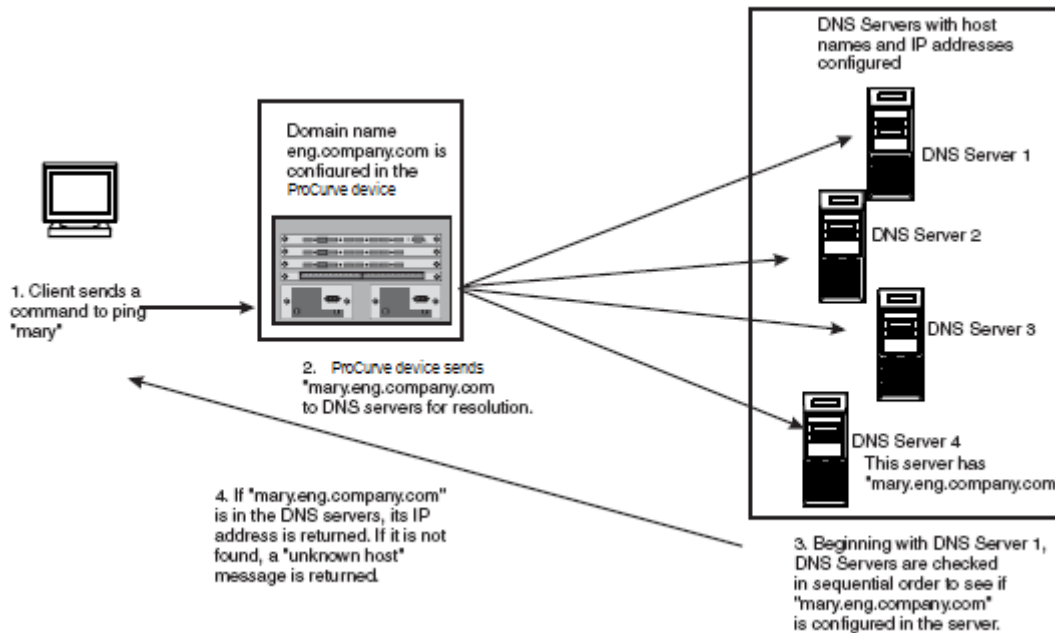


Figure 4 DNS Resolution with One Domain Name

Beginning with software release 08.0.00, you can create a list of domain names that can be used to resolve host names. This list can have more than one domain name. When a client performs a DNS query all hosts within the domains on the list can be recognized and queries can be sent to any domain on the list.

Also, the ProCurve device now contains a DNS cache table that contains a list of host names that have been resolved to their IP addresses. This DNS cache table allows DNS queries to be processed quickly. When a DNS query is made, the query can be sent to the DNS cache table. If a match is found, the DNS query is resolved. If no match is found, the DNS query is sent to the DNS server to be resolved before any action can be taken.

For example, in Figure 5, the client sends a ping to "mary", which is in the eng.company.com domain. The ProCurve device appends the first domain name on the domain name list to "mary" and sends "mary.eng.company.com" to the DNS cache table. Since "mary.eng.company.com" is in the DNS cache table, the query is resolved quickly and the ProCurve device returns the IP address of "mary".

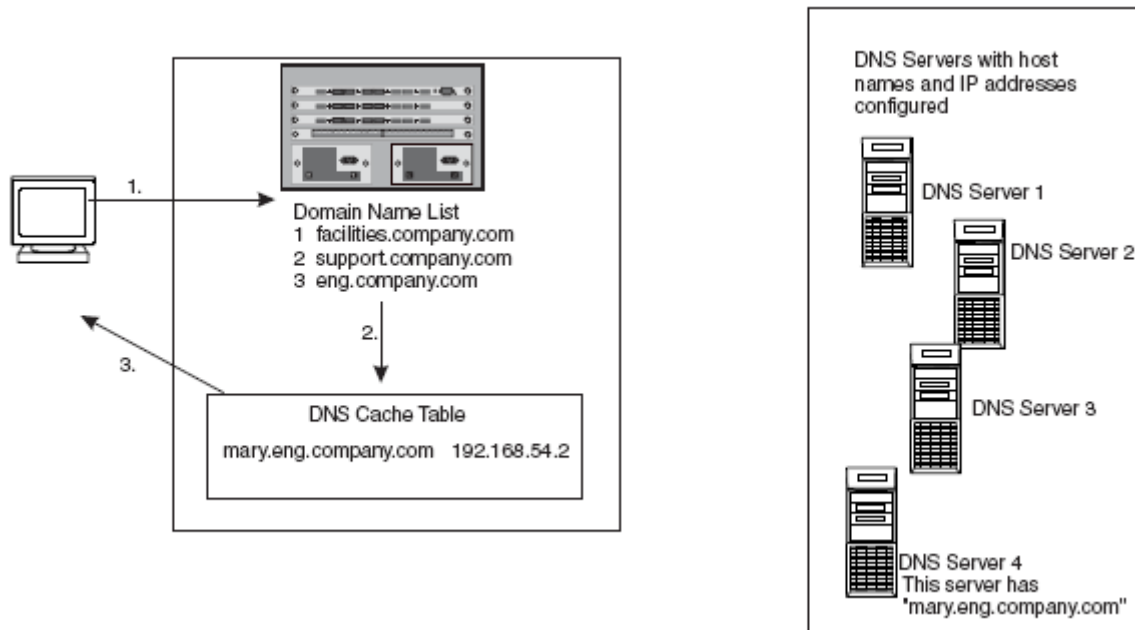


Figure 5 DNS Resolution with Host Name in DNS Cache Table

However, if “mary.eng.company.com” is not in the DNS cache table, as in Figure 6, the host name is resolved as follows:

1. A command to ping “mary” is entered at the client.
2. The ProCurve device appends the first domain name to “mary” and sends the qualified host name “mary.facilities.company.com” to the DNS Cache table.
3. The DNS cache table does not have a “mary.facilities.company.com” entry, so it sends the host name to the DNS servers. Each DNS server is tried in sequential order.
4. Since none of the DNS servers have an entry to “mary.facilities.company.com”, the request is sent back to the Domain Name List. The next domain name is appended to “mary”.
Step 2 through Step 4 are repeated until all the domains in the domain name list are tried. In Figure 6, “mary.eng.company.com” is found in DNS Server 4.
5. Since a match is found, the host name “mary.eng.company.com” and its IP address is added to the DNS cache table.

- The host's IP address is returned to the client. However, if no match is found, an "unknown host" message is returned to the client.

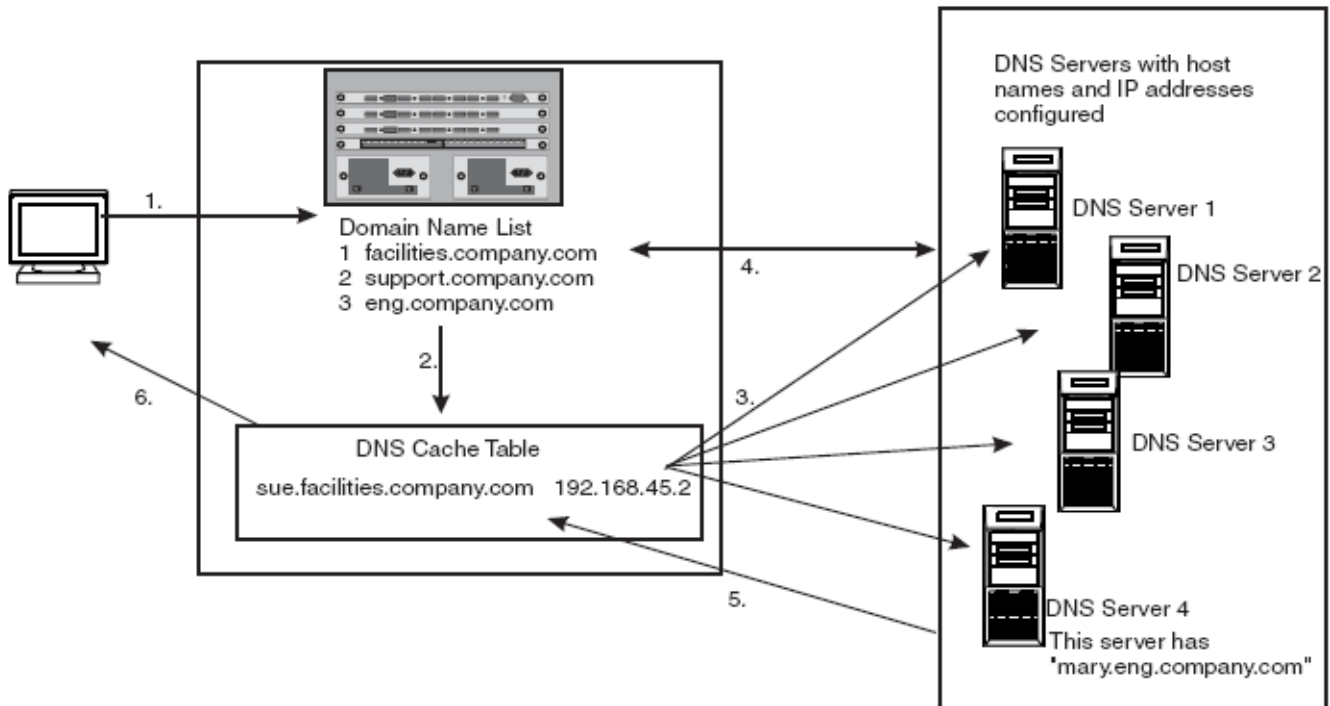


Figure 6 DNS resolution with host name not in DNS Cache Table

Over a period of time, there may be changes to the information in the DNS cache table. For example, a host's IP address can change, making the entries in the DNS cache table invalid. The ProCurve device polls each entry in the DNS cache table to determine if the information in the DNS cache table is still valid. By default, the ProCurve device sends a ping to the host every 1 minute. This polling interval can be changed.

Defining a Domain Name

If you want to define only one domain to resolve host names, enter a command such as the following:

```
ProCurve 9300(config)# ip dns domain-name eng.company.com
```

Syntax: [no] ip dns domain-name <domain-name>

Enter the domain name for <domain-name>.

Defining a Domain List

If you want to use more than one domain name to resolve host names, you can create a list of domain names. For example, enter the commands such as the following:

```
ProCurve 9300(config)# ip dns domain-list facilities.company.com
ProCurve 9300(config)# ip dns domain-list support.company.com
ProCurve 9300(config)# ip dns domain-list eng.company.com
```

The domain names are tried in the order you enter them

Syntax: [no] ip dns domain-list <domain-name>

Enter the full domain name for <domain-name>.

Use the **no** form of the command to remove a domain name.

Displaying the Domain Name List

To determine what domain names have been configured in the domain list, enter the following command:

```
ProCurve 9300(config)#show ip dns domain-list

1 facilities.company.com
2 support.company.com
3 eng.company.com
ProCurve 9300(config)#
```

Syntax: show ip dns domain-list

Defining DNS Servers

You can configure ProCurve device to recognize up to four DNS servers. The first entry serves as the primary (default) DNS server. If a query to the primary DNS server fails to be resolved after three attempts, the next DNS server is queried (also up to three times). This process continues for each defined DNS server until the query is resolved. The order in which the DNS servers are polled is the same as the order in which you enter them.

To define DNS servers, enter a command such as the following:

```
ProCurve 9300(config)# ip dns server-address 209.157.22.199 205.96.7.15 208.95.7.25
201.98.7.15
```

Syntax: [no] ip dns server-address <ip-addr> [<ip-addr>] [<ip-addr>] [<ip-addr>]

In the example above, the first DNS server entered becomes the primary DNS server and all others are secondary servers. Because DNS IP address 201.98.7.15 is the last DNS server listed, it is also the last DNS server consulted to resolve a query.

Verifying Domain Name or IP Address

You can use the **ip domain-lookup** command to verify the host name for an IP address or the IP address for a host name. For example, if you have an IP address and you want to find out what host name it resolves to, enter the following command:

```
ProCurve 9300#ip domain-lookup 66.151.144.5

Host                               Flag      TTL/min  Type  Address
border2.pc0-0-bbnet1.sje.pnap.net  (TMP,OK) 720      IP    66.151.144.5
ProCurve 9300#
```

You can also enter the following:

```
ProCurve 9300#ip domain-lookup border2
Host                               Flag      TTL/min  Type  Address
border2.pc0-0-bbnet1.sje.pnap.net  (TMP,OK) 720      IP    66.151.144.5
ProCurve 9300#
```

Syntax: ip domain-loopkup <ip-address> | <host-name>

Enter an IP address to obtain the host name. Enter the host name to obtain the IP address.

The complete, qualified host name, along with its IP address and TTL value are displayed.

Adding Host Names to the DNS Cache Table

The entries in a DNS cache table are used to resolve host names to IP addresses. When a client initiates a DNS query, the ProCurve device checks the DNS cache table to see if the host name can be resolved to any of the entries. If a match is found, the query is resolved. If a match is not found, the DNS resolver sends the query to the DNS servers. If the name is resolved, the complete, qualified host name and its IP address is added to the DNS cache table and the hosts' IP address is returned to the client.

You can manually add entries to the DNS cache table if you know a host's complete, qualified name and its IP address. To add host names and their IP addresses to the DNS cache table, enter commands such as the following:

```
ProCurve 9300(config)# ip dns cache-entry www.procurve.com 63.236.63.244 720
dynamic-cache-entry
```

Syntax: [no] ip dns cache-entry <host-name> <ip-address> <tll-value> dynamic-cache entry | static-cache-entry

Enter a complete, qualified name for <host-name>. For example, enter www.company.com or host.company.com.

Enter the IP address of the host. This must be the correct IP address for the host.

Enter a time to live (TTL) value for <tll-value>. The TTL determines how many minutes host information stays in the DNS cache table if it has been dynamically added. Once the TTL value expires, the dynamically added host is removed from the table. If the host is added as a static host, the TTL value never changes and the entry does not expire unless it is manually removed from the table.

Enter **dynamic-cache-entry** if you want the host to be listed in the DNS cache table for the duration of the TTL value you entered. Once the TTL value expires, the domain is removed from the DNS cache table.

Enter **static-cache-entry** if you want the domain to remain in the DNS cache table until it is manually cleared.

Use the **no** form of the command to manually remove an entry from the DNS cache table; however, you must enter the entire entry to delete the entry. For example, you must enter:

```
ProCurve 9300(config)# no ip dns cache-entry www.procurve.com 192.6.234.18 720
dynamic-cache-entry
```

Use the **clear ip dns cache-table** command to clear all the entries in the DNS cache table.

Clearing the DNS Cache Table

To clear the entire DNS cache table, enter the following command:

```
ProCurve 9300#clear ip dns cache-table
```

Syntax: clear ip dns cache-table

Displaying the DNS Cache Table

To display what hosts are currently in the DNS cache table, enter the following command:

```
ProCurve 9300(config)#show ip dns cache-table

Host                               Flag      TTL/min  Type  Address
border2.pc0-0-bbnet1.sje.pnap.net  (TMP,OK) 720      IP    66.151.144.5
sl-internap-109-0.sprintlink.net   (TMP,OK) 1440     IP    144.223.242.86
sl-st21-sj-13-0.sprintlink.net     (TMP,OK) 1440     IP    144.232.20.59
sl-bb21-sj-12-0.sprintlink.net     (TMP,OK) 1440     IP    144.232.3.201
sl-bb24-sj-9-0.sprintlink.net      (TMP,OK) 1440     IP    144.232.20.181
sl-bb21-stk-9-0.sprintlink.net     (TMP,OK) 1440     IP    144.232.4.245
sl-gw27-stk-1-2.sprintlink.net     (TMP,OK) 4319     IP    144.228.145.69
core1.ge0-1-bbnet2.sfj.pnap.net    (TMP,OK) 719      IP    216.52.0.65
www.australia.com                  (TMP,OK) 14       IP    66.151.135.102
mail.company.com                   (STA,OK) 26       IP    64.236.22.148
ProCurve 9300(config)#
```

Syntax: show ip dns cache-table

This Column...	Displays...
Host	The complete, qualified domain name of the host.
Flag	Indicates if the entry is dynamic or static and if the information for the domain is up to date: <ul style="list-style-type: none"> • TMP – Entry is dynamic • STA – Entry is static • OK – Information for the entry is up to date • EX – Information for the entry is no longer valid
TTL/min	If the entry is dynamic (TMP) this value shows how long the entry remains in the DNS cache table. If the entry is static (STA), it remains in the DNS cache table and never changes until it is manually removed or the DNS cache table is cleared, even if it shows a TTL/min value.
Type	Type of IP address stored for the entry.
Address	The IP address of the entry.

Defining the Polling Interval

The polling interval determines how often the ProCurve device pings a host in the DNS cache table to determine if the information for that host has changed. If the ping request is successful, an OK value is entered for the host in the DNS cache table. If the host has been added as a dynamic entry, its current TTL value is also updated. If the ping request is unsuccessful (for example, if the host's IP address is no longer valid) an EX value is entered for the host.

To define a polling interval, enter a command such as the following:

```
ProCurve 9300(config)# ip dns poll-interval 7
```

Syntax: ip dns poll-interval <minutes>

Enter the polling interval in minutes. The default is 1 minute.

Displaying the Polling Interval

To display the current polling interval configured for the device, enter the following command:

```
ProCurve 9300(config)#show ip dns poll-time-interval
```

```
Current DNS polling interval is 7 minutes
```

```
ProCurve 9300(config)#
```

Syntax: [no] show ip dns poll-time-interval

Using a DNS Name To Initiate a Trace Route

Suppose you want to trace the route from a ProCurve routing switch to a remote server identified as NYC02 on domain newyork.com. Because the NYC02@ds1.newyork.com domain is already defined on the routing switch, you need to enter only the host name, NYC02, as noted below.

```
ProCurve 9300# traceroute nyc02
```

Syntax: traceroute <host-ip-addr> [maxttl <value>] [minttl <value>] [numeric] [timeout <value>] [source-ip <ip addr>]

The only required parameter is the IP address of the host at the other end of the route. See the *Command Line Interface Reference* for information about the parameters.

After you enter the command, a message indicating that the DNS query is in process and the current DNS address (IP address of the domain name server) being queried appear on the screen:

```
Type Control-c to abort
Sending DNS Query to 209.157.22.199
Tracing Route to IP node 209.157.22.80
To ABORT Trace Route, Please use stop-traceroute command.
Traced route to target IP node 209.157.22.80:
  IP Address          Round Trip Time1    Round Trip Time2
207.95.6.30          93 msec             121 msec
```

NOTE: In the above example, 209.157.22.199 is the IP address of the domain name server (default DNS address), and 209.157.22.80 represents the IP address of the NYC02 host.

IGMP V3 Snooping

Release 08.0.00 adds support for IGMP V3 snooping. This section consists of the following topics:

- “IGMP V3 Snooping Overview” on page 54
- “Configuring IGMP V3 Snooping” on page 55
- “Displaying IGMP V3 Snooping Information” on page 58

IGMP V3 Snooping Overview

This section presents an overview of IGMP V3 snooping on ProCurve devices. For additional information on IGMP V3, see the *Advanced Configuration and Management Guide*.

Using IGMP Protocols on ProCurve Devices

The default behavior for a ProCurve device to handle multicast packets is to broadcast them to every port in the VLAN, except the incoming port (unless the **route-only** command is configured). Internet Group Management Protocol (IGMP) protocols allow the ProCurve device to forward multicast traffic to the ports that want it, and stop forwarding multicast traffic to ports that don't want it.

Clients (hosts) send IGMP membership reports to an IGMP-enabled ProCurve device to indicate the requested traffic stream. The ProCurve device periodically broadcasts IGMP queries to all ports in the VLAN. A client must send membership reports immediately when it first intends to receive multicast traffic. The client also responds to queries from a ProCurve device, and finally sends out a leave message when it no longer wants traffic.

An IGMP-enabled ProCurve device's responsibility is to broadcast general queries periodically, and to send group queries when receiving a leave message to make sure that no other clients on the same port still want this specific traffic before removing traffic from the port.

IGMP V2 lets clients specify which group (destination IP address) to receive. The protocol cannot choose the source of the traffic. In contrast, IGMP V3 is for source-specific multicast traffic. It adds the capability for clients to include or exclude specific traffic sources. An IGMP V3 device's port state could be in include or exclude mode; there are six types of group records for client reports. See “IGMP V3” in the *Advanced Configuration and Management Guide* for detailed IGMP V3 information.

IGMP protocols provide a scheme for clients and a router to exchange messages and let the device build a database indicating which port wants what traffic. IGMP protocols do not specify forwarding methods, however. IGMP snooping or multicast protocols such as PIM or DVMRP are required to handle packet forwarding. PIM and DVMRP can route multicast packets within and outside a VLAN, while IGMP snooping can switch packets only within a VLAN. Servers (traffic sources) are not required to send IGMP memberships. For basic IGMP snooping information, see the “Configuring IP Multicast Traffic Reduction” chapter in the *Installation and Basic Configuration Guide*.

IGMP Snooping Support on ProCurve Devices

ProCurve devices running router software support IGMP V2 snooping using Layer 4 CAM starting with release 07.7.00, as well as IGMP V3 for PIM and DVMRP starting with release 07.8.00. Release 08.0.00 adds support for IGMP V3 snooping.

IGMP V3 is a source-specific protocol, and requires Layer 4 CAM to match both source and group. If you configure IGMP V3 on a VLAN, the VLAN uses Layer 4 CAM.

When Layer 2 CAM is used, traffic is switched solely based on destination MAC address. Consequently, traffic of the same group coming to the same port, regardless of its source, is switched in the same way. In addition, the lowest 23 bits of the group address are mapped to a MAC address. In this way, multiple groups (for example, 224.1.1.1 and 225.1.1.1) have the same MAC address. Groups having the same MAC address are switched to the same destination ports, which are the superset of individual group output ports. Thus, the use of Layer 2 CAM might cause unwanted packets to be sent to some ports. However, the switch generally needs far less Layer 2 CAM than it does Layer 4 CAM, which is required for each stream with a different source and group. As a result, the use of Layer 4 CAM should be avoided if there are many (for example, one thousand) different source and group pairs.

Configuring Queriers and Non-Queriers

An IGMP-enabled ProCurve device can be configured to be a querier (active) or non-querier (passive). An IGMP querier sends queries; a non-querier listens for IGMP queries and forwards them to the entire VLAN. Starting in release 08.0.00, VLANs can be independently configured to be queriers or non-queriers. If a VLAN has a connection to a PIM or DVMRP-enabled port on another router, the VLAN should be configured as a non-querier. When two IGMP snooping devices are connected together, and there is no connection to a PIM or DVMRP-enabled port, then one of the devices should be configured as a querier. If both devices are configured as queriers, then one of them will stop sending queries after the two devices exchange queries. If the querier is a switch, it must have a global IP address in order to send queries. If the querier is a router, the snooping VLAN's router interface must have an IP address, or the router must have a global loopback address.

VLAN-Specific Configuration

Starting in release 08.0.00, you can configure snooping on some VLANs or all VLANs. Each VLAN can independently enable or disable IGMP or PIM snooping, or can be configured with IGMP V2 or V3. In general, the **ip multicast...** commands apply globally to all VLANs except those configured with VLAN-specific **multicast...** commands. The VLAN-specific **multicast...** commands supersede the global **ip multicast...** commands. Per-VLAN configuration is available starting in release 08.0.00.

The configuration of IGMP for snooping and for PIM/DVMRP are independent. The **ip igmp...** commands, available in router code, set IGMP parameters used by PIM/DVMRP, while the **ip multicast...** and **multicast...** commands apply to snooping. In router code, if snooping is configured for a VLAN, and the VLAN's associated virtual interface has PIM or DVMRP enabled, then PIM or DVMRP has higher priority over snooping. The output of the **show multicast vlan** command indicates whether snooping is disabled on a VLAN because PIM or DVMRP is enabled.

Using IGMP V2 with IGMP V3

As with the IGMP V3 functionality introduced in release 07.8.00, snooping can be configured as IGMP V2 or V3 on individual ports on a VLAN.

An interface or router sends the queries and reports that include its IGMP version specified on it. It may recognize a query or report that has a different version, but it may not be able to process them.

Also, a device running IGMP V3 can recognize and process IGMP V2 packets, but when that device sends queries to an IGMP V2 interface, the host on that interface may not recognize the IGMP V3 queries. The interface or router does not automatically downgrade the IGMP version running on them to avoid version deadlock. The **show ip multicast traffic** command displays the number of packets received on a port that have a different IGMP version.

Configuring IGMP V3 Snooping

Configuring IGMP V3 snooping on a ProCurve device consists of the following global and VLAN-specific tasks:

Global tasks:

- Configuring the IGMP mode: active or passive

- Modifying the age interval
- Configuring filtering for multicast groups
- Dropping IGMP V3 traffic in hardware
- Specifying the interval for query messages (active IGMP mode only)
- Specifying that all VLANs use Layer 4 CAM for IGMP snooping
- Specifying the global IGMP version

VLAN specific tasks:

- Configuring the IGMP mode for the VLAN: active or passive
- Enabling or disabling IGMP snooping for the VLAN
- Enabling or disabling PIM SM snooping
- Configuring the IGMP version for the VLAN
- Configuring the IGMP version for individual ports in the VLAN
- Enabling client tracking and the fast-leave feature

Configuring the Global IGMP Mode

You can use active or passive IGMP mode on the ProCurve device. The default mode is passive. If you specify an IGMP version for a VLAN, it overrides the global setting.

- Active – When active IGMP mode is enabled, a ProCurve device actively sends out IGMP queries to identify IP multicast groups on the network and makes entries in the IGMP table based on the Group Membership reports received from the network.
- Passive – When passive IGMP mode is enabled, the device listens for IGMP Group Membership reports but does not send IGMP queries.

To set the IGMP mode for the ProCurve device to active, enter the following command:

```
ProCurve 9300(config)# ip multicast active
```

Syntax: [no] ip multicast [active | passive]

If you omit both the **active** and **passive** keywords, it is equivalent to entering **ip multicast passive**.

Modifying the Age Interval

When the ProCurve device receives a Group Membership report, it makes an entry in the IGMP group table for the group in the report. The age interval specifies how long the entry can remain in the table without the device receiving another Group Membership report.

To modify the age interval, enter a command such as the following:

```
ProCurve 9300(config)# ip multicast age-interval 280
```

Syntax: [no] ip multicast age-interval <interval>

The <interval> parameter specifies the interval between queries. You can specify a value from 10 – 1220 seconds. The default is 140 seconds.

Configuring Filtering for Multicast Groups

By default, ProCurve devices forward multicast traffic for all valid multicast groups. You can configure a ProCurve device to filter out all multicast traffic for groups other than the ones for which the device has received Group Membership reports.

When device starts up, it forwards all multicast groups even though IGMP snooping is configured. This process continues until the device receives a group membership report. Once the group membership report is received, the device drops all multicast packets for groups other than the ones for which the device has received the group membership report. When there are no members in a group, the device drops all multicast packets for that group. If you have multicast applications for which the client never sends a group membership report, you should not enable IP multicast filtering.

To enable IP multicast filtering, enter the following command:

```
ProCurve 9300(config)# ip multicast filter
```

Syntax: [no] ip multicast filter

Dropping IGMP V3 Traffic in Hardware

When there are no clients for a flow, and the **ip multicast filter** command is configured, you can configure the device to drop the IGMP V3 traffic in hardware.

To cause IGMP V3 traffic to be dropped in hardware when there are no clients for a flow, enter the following command:

```
ProCurve 9300(config)# ip multicast hardware-drop
```

Syntax: [no] ip multicast hardware-drop

Modifying the Query Interval (Active IGMP Mode Only)

If the IGMP mode is set to active, you can modify the query interval, which specifies how often a ProCurve device sends Group Membership queries.

To modify the query interval, enter a command such as the following:

```
ProCurve 9300(config)# ip multicast query-interval 120
```

Syntax: [no] ip multicast query-interval <interval>

The <interval> parameter specifies the interval between queries. You can specify a value from 10 – 600 seconds. The default is 60 seconds.

Configuring the Global IGMP Version

You can specify the global IGMP version used on the ProCurve device, either IGMP V2 or IGMP V3. For example, the following command causes the ProCurve device to use IGMP V3:

```
ProCurve 9300(config)# ip multicast version 3
```

Syntax: [no] ip multicast version 2 | 3

In addition, you can optionally specify the IGMP version for individual VLANs, or individual ports within VLANs. If no IGMP version is specified for a VLAN, then the globally configured IGMP version is used. If an IGMP version is specified for individual ports in a VLAN, those ports use that version, instead of the version specified for the VLAN or the globally specified version.

Configuring the IGMP Mode for a VLAN

You can use active or passive IGMP mode on a VLAN. The default mode is passive. The setting specified for the VLAN overrides the global setting.

- Active – When active IGMP mode is enabled, a ProCurve device actively sends out IGMP queries to identify IP multicast groups on the network and makes entries in the IGMP table based on the Group Membership reports received from the network.
- Passive – When passive IGMP mode is enabled, the device listens for IGMP Group Membership reports but does not send IGMP queries.

To set the IGMP mode for VLAN 20 to active, enter the following commands:

```
ProCurve 9300(config)# vlan 20
ProCurve 9300(config-vlan-20)# multicast active
```

Syntax: [no] multicast active | passive

Disabling IGMP Snooping for the VLAN

When IGMP snooping is enabled globally, you can disable it for a specific VLAN. For example, the following commands cause IGMP snooping to be disabled for VLAN 20. This setting overrides the global setting.

```
ProCurve 9300(config)# vlan 20
ProCurve 9300(config-vlan-20)# disable-igmp-snoop
```

Syntax: [no] disable-igmp-snoop

Disabling PIM SM Snooping for the VLAN

When PIM SM snooping is enabled globally, you can disable it for a specific VLAN. For example, the following commands cause PIM SM snooping to be disabled for VLAN 20. This setting overrides the global setting.

```
ProCurve 9300(config)# vlan 20
ProCurve 9300(config-vlan-20)# disable-pimsm-snoop
```

Syntax: [no] disable-pimsm-snoop

Configuring the IGMP Version for the VLAN

You can specify the IGMP version for the ports in a VLAN. For example, the following commands cause VLAN 20 to use IGMP V3:

```
ProCurve 9300(config)# vlan 20
ProCurve 9300(config-vlan-20)# multicast version 3
```

Syntax: [no] multicast version 2 | 3

If no IGMP version is specified, then the globally configured IGMP version is used. If an IGMP version is specified for individual ports in the VLAN, those ports use that version, instead of the version specified for the VLAN.

Configuring the IGMP Version for Individual Ports in the VLAN

You can specify the IGMP version for individual ports in a VLAN. For example, the following commands cause ports 4/3, 4/5, 4/6, and 4/7 to use IGMP V3. The other ports in the VLAN use the IGMP version specified with the **multicast version** command, or if the **multicast version** command is not configured, the globally configured IGMP version is used.

```
ProCurve 9300(config)# vlan 20
ProCurve 9300(config-vlan-20)# multicast port-version 3 ethe 4/3 ethe 4/5 to 4/7
```

Syntax: [no] multicast port-version 2 | 3 <port-numbers>

Enabling Membership Tracking and Fast Leave for the VLAN

IGMP V3 provides membership tracking and fast leave to clients. In IGMP V2, only one client on an interface needs to respond to a router's queries; therefore, some of the clients may be invisible to the router, making it impossible for the router to track the membership of all clients in a group. Also, when a client leaves the group, the router sends group specific queries to the interface to see if other clients on that interface need the data stream of the client who is leaving. If no client responds, the router waits three seconds before it stops the traffic.

When the tracking feature is enabled, the device immediately stops forwarding multicast traffic to the interface (without waiting three seconds) if an IGMP V3 client sends a leave message and there are no other clients. If a client does not send a report during the specified group membership time (the default is 140 seconds), that client is removed from the tracking list.

To enable the tracking and fast leave feature for VLAN 20, enter the following commands:

```
ProCurve 9300(config)# vlan 20
ProCurve 9300(config-vlan-20)# multicast tracking
```

Syntax: [no] multicast tracking

The membership tracking and fast leave features are supported for IGMP V3 only. If any port or any client is not configured for IGMP V3, then the **multicast tracking** command is ignored.

Displaying IGMP V3 Snooping Information

You can display the following IGMP V3 snooping information:

- IGMP error information
- Information about VLANs using Layer 4 CAM
- Hardware resource usage for VLANs using Layer 2 CAM
- Group and forwarding information for VLANs using Layer 2 CAM
- Multicast forwarding cache information

- PIM SM snooping information for VLANs using Layer 2 CAM
- IGMP V3 memory pool usage
- Status of IGMP V3 traffic
- IGMP V3 information by VLAN

Displaying IGMP Error Information

To display information about possible IGMP errors, enter the following command:

```
ProCurve 9300# show ip multicast error
**** Warning! counter igmp checksum error = 10
**** Warning! counter igmp, pkt buf alloc fail = 7
**** Warning! counter snoop router fid alloc fail = 12
```

Syntax: show ip multicast error

Displaying Information about VLANs Using Layer 4 CAM

You can display the status of all or specific groups of VLANs using Layer 4 CAM by entering the following command:

```
ProCurve 9300# show ip multicast group
VL20 : 1 groups, 1 group-port
      group          phy-port static querier life mode   #_src
1     224.1.1.1      e4/12  no    no    140  exclude 0
```

To display detailed information, enter the following command:

```
ProCurve 9300# show ip multicast group 224.1.1.1 detail
Display group 224.1.1.1 in all interfaces in details.
VL20 : 1 groups, 1 group-port
      group          phy-port static querier life mode   #_src
1     224.1.1.1      e4/12  no    no    include 2
      group: 224.1.1.1, include, permit 2 (source, life):
      (1.1.32.1 120) (1.1.32.2 120)
```

If the tracking and fast leave features are enabled, you can display the list of clients that belong to a particular group by entering the following command:

```
ProCurve 9300# show ip multicast group 224.1.1.1 track
Display group 224.1.1.1 in all interfaces with tracking enabled.
VL20 : 1 groups, 1 group-port, tracking_enabled
      group          phy-port static querier life mode   #_src
1     224.1.1.1      e4/12  no    no    include 2
      receive reports from 1 clients:
      2.2.2.100
```

Syntax: show ip multicast group [<group-address>] [detail] [tracking]

If you want a report for a specific multicast group, enter that group's address for <group-address>. Omit the <group-address> if you want a report for all multicast groups using Layer 4 CAM.

Enter **detail** if you want to display the source list of the multicast group.

Enter **tracking** if you want information on interfaces that have tracking enabled.

Table 11 describes the information displayed by the **show ip multicast group** command.

Table 11: Output from the show ip multicast group command

This Field	Displays
Group	The address of the multicast group (destination IP address)
Phy-port	The physical port on which the multicast group was received.
Static	A "yes" entry in this column indicates that the multicast group was configured as a static group; "No" means it was not. Since static groups can be configured only for PIM/DVMRP, this should always be "no" for snooping.
Querier	"Yes" means that the port is a querier port; "No" means it is not. A port becomes a non-querier port when it receives a query from a source with a lower source IP address than the port.
Life	Shows the number of seconds the interface can remain in exclude mode. An exclude mode changes to include mode if it does not receive an "IS_EX" or "TO_EX" message during a certain period of time. The default is 140 seconds. There is no "life" displayed in include mode.
Mode	Indicates current mode of the interface: Include or Exclude. If the interface is in Include mode, it admits traffic only from the source list. If an interface is in Exclude mode, it denies traffic from the source list and accepts the rest.
#_src	Identifies the source list that will be included or excluded on the interface. If an IGMP V2 group is in Exclude mode with a #_src of 0, the group excludes traffic from 0 (zero) source list, which means that all traffic sources are included.
Group	If you requested a <i>detailed</i> report, the following information is displayed: <ul style="list-style-type: none"> • The multicast group address • The mode of the group • A list of sources from which traffic will be admitted (include) or denied (exclude) on the interface is listed. • The life of each source list. If you requested a <i>tracking</i> report, the clients from which reports were received are identified.

Displaying Multicast Forwarding Cache Information

The multicast forwarding cache contains multicast information for VLANs using Layer 4 CAM. To display information in the multicast forwarding cache, enter the following command:

```
ProCurve 9300# show ip multicast mcache
mcache is only available for vlans using L4 cam
Example: (S G) in tag e4/3 cnt=: e4/3 is input, cnt: SW proc. count <--- new
      OIF: TR(e4/1, e4/2) (20): 4/1 is primary trunk, 4/2 is real out, (20):
age <--- new
vlan 1, has 0 cache
vlan 20, has 3 cache
1 (1.1.32.3 224.1.1.1) in tag e4/3, cnt=1
  OIF: tag TR(e4/1,e4/1) (20), e4/7 (20),
  age=0s up-time=0m fid=2d7f, cam=00aca,
2 (1.1.32.2 224.1.1.1) in tag e4/3, cnt=1
  OIF: tag TR(e4/1,e4/2) (20), e4/7 (20),
  age=0s up-time=0m fid=2d43, cam=00be1,
3 (1.1.32.0 224.1.1.1) in tag e4/3, cnt=2
  OIF: tag TR(e4/1,e4/1) (20), e4/7 (20),
  age=0s up-time=0m fid=08b5, cam=00aa0,
```

Syntax: show ip multicast mcache

Table 12 describes the information displayed by the **show ip multicast mcache** command.

Table 12: Output from the show ip multicast mcache command

This Field...	Displays...
tag	Whether the port is tagged.
cnt	The number of packets processed in software. This does not include hardware-switched packets.
OIF	The output interface.
TR(e4/1, e4/2)	Trunk port information. In this example, e4/1 is the primary port, e4/2 is the real output port.
(20)	Amount of time since receiving IGMP membership for this port.

To clear the multicast forwarding cache of information about VLANs using Layer 4 CAM, enter the following command:

```
ProCurve 9300# clear ip multicast mcache
```

Syntax: clear ip multicast mcache

To clear the multicast forwarding cache of information about a specific VLAN that uses Layer 4 CAM, enter a command such as the following:

```
ProCurve 9300# clear ip multicast vlan 20
```

Syntax: clear ip multicast vlan <vlan-id>

When a VLAN uses Layer 4 CAM, all PIM sparse join and prune messages are directly added to the multicast forwarding cache; to display PIM sparse snooping information for VLANs using Layer 4 CAM, use the **show ip multicast mcache** command.

Displaying IGMP V3 Memory Pool Usage

You can display information about the memory pools used for IGMP. ProCurve devices have a single set of memory pools for Layer 4 CAM.

Displaying Status of IGMP V3 Traffic

To display status information for IGMP V3 traffic, enter the following command:

```
ProCurve 9300# show ip multicast traffic
Total Recv: 42829, Xmit: 0 (including IGMP for PIM/DVMRP)
Recv  QryV2  QryV3  G-Qry  GSQry  MbrV2  MbrV3  Leave  IsIN  IsEX  ToIN  ToEX  ALLO  BLK
VL1   0      0      0      0      0      0      0      0      0      0      0      0      0
VL20  193  34551  1      3      7704   12     365    9     7704  365    0      0      3
Send  QryV1  QryV2  QryV3  G-Qry  GSQry
VL1   0      0      0      0      0
VL20  0      0      0      0      0
VL1   pimsm-snooping, Hello: 0, Join/Prune: 0
VL20  pimsm-snooping, Hello: 44683, Join/Prune: 32849
```

Syntax: show ip multicast traffic

Table 13 describes the information displayed by the **show ip multicast traffic** command.

Table 13: Output from the show ip multicast traffic command

This Field	Displays
QryV2	Number of general IGMP V2 queries received or sent by the virtual routing interface.
QryV3	Number of general IGMP V3 queries received or sent by the virtual routing interface.
G-Qry	Number of group specific queries received or sent by the virtual routing interface.
GSQry	Number of source specific queries received or sent by the virtual routing interface.
MbrV2	The IGMP V2 membership report.
MbrV3	The IGMP V3 membership report.
Leave	Number of IGMP V2 "leave" messages on the interface. (See ToEx for IGMP V3.)
IsIN	Number of source addresses that were included in the traffic.
IsEX	Number of source addresses that were excluded in the traffic.
ToIN	Number of times the interface mode changed from exclude to include.
ToEX	Number of times the interface mode changed from include to exclude.
ALLO	Number of times that additional source addresses were allowed or denied on the interface.
BLK	Number of times that sources were removed from an interface.

To clear the counters for IGMP V3 traffic, enter the following command:

```
ProCurve 9300# clear ip multicast traffic
```

Syntax: clear ip multicast traffic

Displaying IGMP V3 Information by VLAN

You can display IGMP V3 information for all VLANs or for a specific VLAN. For example, to display multicast information for VLAN 20, enter the following command:

```
ProCurve 9300# show ip multicast vlan 20
version=2, query-t=20, group-aging-t=140, max-resp-t=3
VL20: dft V2, L4 CAM, glb cfg Passive, pimsm (glb cfg), 0 grp, 4 caches, , rtr-
fid=08A6
  router ports: e4/3(160), e4/7(180), e4/1(140),
e4/7   has    0 groups, non-QR (passive), default V2
e4/3   has    0 groups, non-QR (passive), default V2
e4/1   has    0 groups, non-QR (passive), default V2
```

Syntax: show ip multicast vlan [<vlan-id>]

If you do not specify a <vlan-id>, information for all VLANs is displayed.

Table 14 describes information displayed by the **show ip multicast vlan** command.

Table 14: Output from the show ip multicast vlan command

This Field	Displays
version	The IGMP version number
query-t	How often a querier sends a general query on the interface.
group-aging-t	The number of seconds multicast groups can be members of this group before aging out.
rtr-fid	The FID of the ports receiving queries
max-resp-t	The maximum number of seconds a client can wait before it replies to the query.
non-QR	Indicates that the port is a non-querier.
QR	Indicates that the port is a querier.

Increased Size of the MSDP Source Active Cache

When a ProCurve device is serving as an MSDP router and RP (Rendezvous Point), then the device can be configured to cache the Source Active messages it receives. In this case, even if the RP does not have a receiver for a group when the RP receives the Source Active message for the group, the RP can immediately send a Join for a new receiver that wants to join the group, without waiting for the next Source Active message from the RP in the source's domain.

Starting in release 08.0.00, the size of the cache used to store MSDP Source Active messages on the ProCurve device has been increased from 4k to 8k.

Specifying a Designated Router Election Priority for PIM V2

In a multi-access network, where two or more routers' interfaces are connected together, it is possible that a packet could reach more than one router. In this kind of network, PIM elects a designated router (DR), whose duties include sending registration or join/prune messages for hosts and forwarding traffic for the LAN segment.

In releases prior to 08.0.00, the router with the highest IP address was elected the DR. Starting in release 08.0.00, you can optionally assign each router a DR election priority. The router with the highest DR election priority is elected the DR. If two or more routers all have the same DR election priority, then the router with the highest IP address is elected the DR. If any of the routers in the multi-access network do not support the DR election priority feature, then the router with the highest IP address is elected the DR.

The DR election priority is set on a per-interface basis. For example, to set a DR election priority of 99 for VE 20, enter the following commands:

```
ProCurve 9300(config)# int ve 20
ProCurve 9300(config-vif-20)# ip pim dr-priority 99
ProCurve 9300(config-vif-20)# exit
```

Syntax: [no] ip pim dr-priority <num>

The priority specified by <num> can be from 0 – 4294967295. The default is 1. Starting in this release, the PIM hello message sent by ProCurve devices always contains the router's DR election priority (either the specified priority or 1 if no priority was specified).

To prevent an interface from being elected a DR, you can either set its priority to 0, or you can increase the priority of all interfaces of other routers in the multi-access network to a number greater than 1.

The output of the **show ip pim interface** command has been enhanced to display the device's DR election priority. For example:

```
ProCurve 9300# show ip pim interface
Interface v20
PIM Sparse
TTL Thres: 1, Enabled, dr-priority=99, DR: itself
Local Address: 1.1.20.4
Neighbor:
  1.1.20.5
```

Support for Standard Multicast MIBs

Support for the following standard multicast MIBs has been added to this release:

- RFC 2932: IPv4 Multicast Routing MIB
- RFC 2933: IGMP MIB
- RFC 2934: PIM for IPv4

The sections below show which objects in the MIBs are supported or not supported, and any limitations regarding their support.

RFC 2932 IPv4 Multicast Routing MIB

Object	Supported?	Object Identifier
ipMRouteEnable	Yes	1.3.6.1.2.1.83.1.1.1
ipMRouteEntryCount	Yes	1.3.6.1.2.1.83.1.1.2.1

ipMRouteTable (IP multicast route table)

ipMRouteGroup	Yes	1.3.6.1.2.1.83.1.1.2.1.1
ipMRouteSource	Yes. Returns the IP address of the multicast server.	1.3.6.1.2.1.83.1.1.2.1.2
ipMRouteSourceMask	Yes. Always shows "1".	1.3.6.1.2.1.83.1.1.2.1.3

Object	Supported?	Object Identifier
ipMRouteUpstreamNeighbor	Yes	1.3.6.1.2.1.83.1.1.2.1.4
ipMRouteInIfIndex	Yes	1.3.6.1.2.1.83.1.1.2.1.5
ipMRouteUpTime	Yes	1.3.6.1.2.1.83.1.1.2.1.6
ipMRouteExpiryTime	Yes	1.3.6.1.2.1.83.1.1.2.1.7
ipMRoutePkts	No	1.3.6.1.2.1.83.1.1.2.1.8
ipMRouteDifferentInIfPackets	Yes	1.3.6.1.2.1.83.1.1.2.1.9
ipMRouteOctets	No	1.3.6.1.2.1.83.1.1.2.1.10
ipMRouteProtocol	Yes	1.3.6.1.2.1.83.1.1.2.1.11
ipMRouteRtProto	Yes	1.3.6.1.2.1.83.1.1.2.1.12
ipMRouteRtAddress	Yes. Returns the IP address of the multicast server.	1.3.6.1.2.1.83.1.1.2.1.13
ipMRouteRtMask	Yes. Always shows "1".	1.3.6.1.2.1.83.1.1.2.1.14
ipMRouteRtType	Yes	1.3.6.1.2.1.83.1.1.2.1.15
ipMRouteHCOctets	No. Always shows "0".	1.3.6.1.2.1.83.1.1.2.1.16

ipMRouteNextHopTable (IP multicast next hop table)

"Next hop" in this table refers to downstream traffic.

ipMRouteNextHopGroup	Yes	1.3.6.1.2.1.83.1.1.3.1.1
ipMRouteNextHopSource	Yes	1.3.6.1.2.1.83.1.1.3.1.2
ipMRouteNextHopSourceMask	Yes	1.3.6.1.2.1.83.1.1.3.1.3
ipMRouteNextHopIfIndex	Yes	1.3.6.1.2.1.83.1.1.3.1.4
ipMRouteNextHopAddress	Yes	1.3.6.1.2.1.83.1.1.3.1.5
ipMRouteNextHopState	Yes. Always shows forwarding(2).	1.3.6.1.2.1.83.1.1.3.1.6
ipMRouteNextHopUpTime	No. Always shows "0".	1.3.6.1.2.1.83.1.1.3.1.7
ipMRouteNextHopExpiryTime	No. Always shows "0".	1.3.6.1.2.1.83.1.1.3.1.8
ipMRouteNextHopClosestMemberHop	No. Always shows "0".	1.3.6.1.2.1.83.1.1.3.1.9
ipMRouteNextHopProtocol	Yes	1.3.6.1.2.1.83.1.1.3.1.10
ipMRouteNextHopPkts	No. Always shows "0".	1.3.6.1.2.1.83.1.1.3.1.11

ipMRouteInterfaceTable (IP multicast route table for interfaces)

ipMRouteInterfaceIfIndex	Yes	1.3.6.1.2.1.83.1.1.4.1.1
ipMRouteInterfaceTtl	Yes. Range: 1—31	1.3.6.1.2.1.83.1.1.4.1.2
ipMRouteInterfaceProtocol	Yes	1.3.6.1.2.1.83.1.1.4.1.3

Object	Supported?	Object Identifier
ipMRouteInterfaceRateLimit	No	1.3.6.1.2.1.83.1.1.4.1.4
ipMRouteInterfaceInMcastOctets	Yes. Returns packet count	1.3.6.1.2.1.83.1.1.4.1.5
ipMRouteInterfaceOutMcastOctets	Yes. Returns packet count	1.3.6.1.2.1.83.1.1.4.1.6
ipMRouteInterfaceHCInMcastOctets	Yes. Returns packet count	1.3.6.1.2.1.83.1.1.4.1.7
ipMRouteInterfaceHCOutMcastOctets	Yes. Returns packet count	1.3.6.1.2.1.83.1.1.4.1.8

IP Multicast Scope Boundary Table (IP multicast scope boundary table)

IpMRouteBoundaryIfIndex	Yes	1.3.6.1.2.1.83.1.1.5.1.1
IpMRouteBoundaryAddress	Yes. Value is obtained from ACLs.	1.3.6.1.2.1.83.1.1.5.1.2
IpMRouteBoundaryAddressMask	Yes. Value is obtained from ACLs.	1.3.6.1.2.1.83.1.1.5.1.3
IpMRouteBoundaryStatus	Yes. Read only.	1.3.6.1.2.1.83.1.1.5.1.4

ipMRouteScopeNameTable (IP multicast scope group name table)

Objects in this table are supported.

RFC 2933 IGMP MIB

Object	Supported?	Object Identifier
--------	------------	-------------------

igmpInterfaceTable (IGMP Interface Table)

igmpInterfaceIfIndex	Yes	1.3.6.1.2.1.85.1.1.1.1
igmpInterfaceQueryInterval	Yes. Global value only.	1.3.6.1.2.1.85.1.1.1.2
igmpInterfaceStatus	Yes	1.3.6.1.2.1.85.1.1.1.3
igmpInterfaceVersion	Yes	1.3.6.1.2.1.85.1.1.1.4
igmpInterfaceQuerier	Yes	1.3.6.1.2.1.85.1.1.1.5
igmpInterfaceQueryMaxResponseTime	Yes. Global value only.	1.3.6.1.2.1.85.1.1.1.6
igmpInterfaceQuerierUpTime	Yes	1.3.6.1.2.1.85.1.1.1.7
igmpInterfaceQuerierExpiryTime	Yes	1.3.6.1.2.1.85.1.1.1.8
igmpInterfaceVersion1QuerierTimer	Yes, but only the following values are supported: <ul style="list-style-type: none"> 0 = no V1 querier 1 = no time 	1.3.6.1.2.1.85.1.1.1.9
igmpInterfaceWrongVersionQueries	Yes	1.3.6.1.2.1.85.1.1.1.10
igmpInterfaceJoins	Yes	1.3.6.1.2.1.85.1.1.1.11

Object	Supported?	Object Identifier
igmpInterfaceProxyIfIndex	No	1.3.6.1.2.1.85.1.1.1.12
igmpInterfaceGroups	Yes	1.3.6.1.2.1.85.1.1.1.13
igmpInterfaceRobustness	Yes. Global value only.	1.3.6.1.2.1.85.1.1.1.14
igmpInterfaceLastMemQueryIntvl	Yes	1.3.6.1.2.1.85.1.1.1.15

igmpCacheTable (IGMP Cache Table)

igmpCacheAddress	Yes	1.3.6.1.2.1.85.1.2.1.1
igmpCacheIfIndex	Yes	1.3.6.1.2.1.85.1.2.1.2
igmpCacheSelf	Yes	1.3.6.1.2.1.85.1.2.1.3
igmpCacheLastReporter	Yes	1.3.6.1.2.1.85.1.2.1.4
igmpCacheUpTime	Yes	1.3.6.1.2.1.85.1.2.1.5
igmpCacheExpiryTime	Yes	1.3.6.1.2.1.85.1.2.1.6
igmpCacheStatus	Yes	1.3.6.1.2.1.85.1.2.1.7
igmpCacheVersion1HostTimer	Yes	1.3.6.1.2.1.85.1.2.1.8

RFC 2934 PIM MIB for IPv4

Object	Supported?	Object Identifier
pimJoinPruneInterval	Yes	1.3.6.1.3.61.1.1.1

pimInterfaceTable (PIM Interface Table)

pimInterfaceIfIndex	Yes	1.3.6.1.3.61.1.1.2.1.1
pimInterfaceAddress	Yes	1.3.6.1.3.61.1.1.2.1.2
pimInterfaceNetMask	Yes	1.3.6.1.3.61.1.1.2.1.3
pimInterfaceMode	Yes	1.3.6.1.3.61.1.1.2.1.4
pimInterfaceDR	Yes	1.3.6.1.3.61.1.1.2.1.5
pimInterfaceHelloInterval	Yes. Global value only.	1.3.6.1.3.61.1.1.2.1.6
pimInterfaceStatus	Yes	1.3.6.1.3.61.1.1.2.1.7
pimInterfaceJoinPruneInterval	Yes. Global value only.	1.3.6.1.3.61.1.1.2.1.8
pimInterfaceCBSRPreference	Yes. Global value only.	1.3.6.1.3.61.1.1.2.1.9

Object	Supported?	Object Identifier
--------	------------	-------------------

pimNeighborTable (PIM Neighbor Table)

pimNeighborAddress	Yes	1.3.6.1.3.61.1.1.3.1.1
pimNeighborIfIndex	Yes	1.3.6.1.3.61.1.1.3.1.2
pimNeighborUpTime	Yes	1.3.6.1.3.61.1.1.3.1.3
pimNeighborExpiryTime	Yes	1.3.6.1.3.61.1.1.3.1.4
pimNeighborMode	Yes	1.3.6.1.3.61.1.1.3.1.5

pimIpMRouteTable (PIM IP Multicast Route Table)

pimIpMRouteUpstreamAssertTimer	Yes	1.3.6.1.3.61.1.1.4.1.1
pimIpMRouteAssertMetric	Yes	1.3.6.1.3.61.1.1.4.1.2
pimIpMRouteAssertMetricPref	Yes	1.3.6.1.3.61.1.1.4.1.3
pimIpMRouteAssertRPTBit	Yes	1.3.6.1.3.61.1.1.4.1.4
pimIpMRouteFlags	Yes	1.3.6.1.3.61.1.1.4.1.5

pimIpMRouteNextHopTable (PIM Next Hop Table)

The pimIpMRouteNextHopTable is not supported.

pimRpTable (PIM RP Table)

pimRpGroupAddress	Yes, but read-only and only active groups.	1.3.6.1.3.61.1.1.5.1.1
pimRpAddress	Yes, but read-only.	1.3.6.1.3.61.1.1.5.1.2
pimRpState	Yes, but read-only and value is always up(1).	1.3.6.1.3.61.1.1.5.1.3
pimRpStateTimer	No	1.3.6.1.3.61.1.1.5.1.4
pimRpLastChange (No	1.3.6.1.3.61.1.1.5.1.5
pimRpRowStatus	Yes, but read-only.	1.3.6.1.3.61.1.1.5.1.6

Object	Supported?	Object Identifier
--------	------------	-------------------

pimRpSetTable (PIM RP Set Table)

pimRpSetGroupAddress	Yes	1.3.6.1.3.61.1.1.6.1.1
pimRpSetGroupMask	Yes	1.3.6.1.3.61.1.1.6.1.2
pimRpSetAddress	Yes	1.3.6.1.3.61.1.1.6.1.3
pimRpSetHoldTime	Yes	1.3.6.1.3.61.1.1.6.1.4
pimRpSetExpiryTime	Yes	1.3.6.1.3.61.1.1.6.1.5
pimRpSetComponent	No	1.3.6.1.3.61.1.1.6.1.6

pimCandidateRPTable (PIM Candidate-RP Table)

pimCandidateRPGroupAddress	Yes	1.3.6.1.3.61.1.1.11.1.1
pimCandidateRPGroupMask	Yes	1.3.6.1.3.61.1.1.11.1.2
pimCandidateRPAddress	Yes	1.3.6.1.3.61.1.1.11.1.3
pimCandidateRPRowStatus	Yes	1.3.6.1.3.61.1.1.11.1.4

pimComponentTable (PIM Component Table)

SET operation for this table is not available, since the BSR is in one domain only. This table has only one row.

Use the CLI command **ip pim border** at the interface level to stop the flooding of the bootstrap messages.

pimComponentIndex	Yes	1.3.6.1.3.61.1.1.12.1.1
pimComponentBSRAddress	Yes	1.3.6.1.3.61.1.1.12.1.2
pimComponentBSRExpiryTime	Yes	1.3.6.1.3.61.1.1.12.1.3
pimComponentCRPHoldTime	Yes	1.3.6.1.3.61.1.1.12.1.4
pimComponentStatus	Yes, but read-only	1.3.6.1.3.61.1.1.12.1.5

Using Multi-Device Port Authentication and 802.1X Security on the Same Port

Starting in release 08.0.00, you can configure the ProCurve device to use multi-device port authentication and 802.1X security on the same port.

- The multi-device port authentication feature allows you to configure a ProCurve device to forward or block traffic from a MAC address based on information received from a RADIUS server. Incoming traffic originating from a given MAC address is switched or forwarded by the device only if the source MAC address is successfully authenticated by a RADIUS server. The MAC address itself is used as the username and password for RADIUS authentication. A connecting user does not need to provide a specific username and password to gain access to the network.
- The IEEE 802.1X standard is a means for authenticating devices attached to LAN ports. Using 802.1X port security, you can configure a ProCurve device to grant access to a port based on information supplied by a client to an authentication server.

For information on configuring the multi-device port authentication feature and 802.1X security on ProCurve devices, see the June 2005 or later version of the *Security Guide*.

When both of these features are enabled on the same port, multi-device port authentication is performed prior to 802.1X authentication. If multi-device port authentication is successful, 802.1X authentication may be performed, based on the configuration of a vendor-specific attribute (VSA) in the profile for the MAC address on the RADIUS server.

When both features are configured on a port, a device connected to the port is authenticated as follows:

1. Multi-device port authentication is performed on the device to authenticate the device's MAC address.
2. If multi-device port authentication is successful for the device, then the ProCurve device checks whether the RADIUS server included the Foundry-802_1x-enable VSA (described in Table 15) in the Access-Accept message that authenticated the device.
3. If the Foundry-802_1x-enable VSA is not present in the Access-Accept message, or is present and set to 1, then 802.1X authentication is performed for the device.
4. If the Foundry-802_1x-enable VSA is present in the Access-Accept message, and is set to 0, then 802.1X authentication is skipped. The device is authenticated, and any dynamic VLANs specified in the Access-Accept message returned during multi-device port authentication are applied to the port.
5. If 802.1X authentication is performed on the device, and is successful, then dynamic VLANs or ACLs specified in the Access-Accept message returned during 802.1X authentication are applied to the port.

If multi-device port authentication fails for a device, then by default traffic from the device is either blocked in hardware, or the device is placed in a restricted VLAN. You can optionally configure the ProCurve device to perform 802.1X authentication on a device when it fails multi-device port authentication. See "Example 2" on page 73 for a sample configuration where this is used.

Configuring Vendor-Specific Attributes on the RADIUS Server

If the RADIUS authentication process is successful, the RADIUS server sends an Access-Accept message to the ProCurve device, authenticating the device. The Access-Accept message can include Vendor-Specific Attributes (VSAs) that specify additional information about the device. If you are configuring multi-device port authentication and 802.1X authentication on the same port, then you can configure the Foundry VSAs listed in Table 15 on the RADIUS server.

Add these Foundry vendor-specific attributes to your RADIUS server's configuration, and configure the attributes in the individual or group profiles of the devices that will be authenticated. Foundry's Vendor-ID is 1991, with Vendor-Type 1.

Table 15: foundry vendor-specific attributes for RADIUS

Attribute Name	Attribute ID	Data Type	Description
Foundry-802_1x-enable	6	integer	<p>Specifies whether 802.1X authentication is performed when multi-device port authentication is successful for a device. This attribute can be set to one of the following:</p> <p>0 Do not perform 802.1X authentication on a device that passes multi-device port authentication. Set the attribute to zero for devices that do not support 802.1X authentication.</p> <p>1 Perform 802.1X authentication when a device passes multi-device port authentication. Set the attribute to one for devices that support 802.1X authentication.</p>

Table 15: foundry vendor-specific attributes for RADIUS

Attribute Name	Attribute ID	Data Type	Description
	7	integer	<p>Specifies whether the RADIUS record is valid only for multi-device port authentication, or for both multi-device port authentication and 802.1X authentication.</p> <p>This attribute can be set to one of the following:</p> <ul style="list-style-type: none"> 0 The RADIUS record is valid only for multi-device port authentication. Set this attribute to zero to prevent a user from using their MAC address as username and password for 802.1X authentication 1 The RADIUS record is valid for both multi-device port authentication and 802.1X authentication.

If neither of these VSAs exist in a device's profile on the RADIUS server, then by default the device is subject to multi-device port authentication (if configured), then 802.1X authentication (if configured). The RADIUS record can be used for both multi-device port authentication and 802.1X authentication.

Sample Configurations

The following are two example of configurations that use multi-device port authentication and 802.1X authentication on the same port.

Example 1

Figure 7 illustrates a sample configuration that uses multi-device port authentication and 802.1X authentication on the same port. In this configuration, a PC and an IP phone are connected to port 1/3 on a ProCurve device. Port 1/3 is configured as a dual-mode port.

The PC transmits untagged traffic, and the IP phone is configured to transmit tagged traffic (VLAN named "IP-Phone-VLAN"). The profile for the PC's MAC address on the RADIUS server specifies that the PC should be dynamically assigned to VLAN "Login-VLAN", and the profile for the IP phone specifies that it should be

dynamically assigned to the VLAN named "IP-Phone-VLAN". When User 1 is successfully authenticated using 802.1X authentication, the PC is then placed in the VLAN named "User-VLAN".

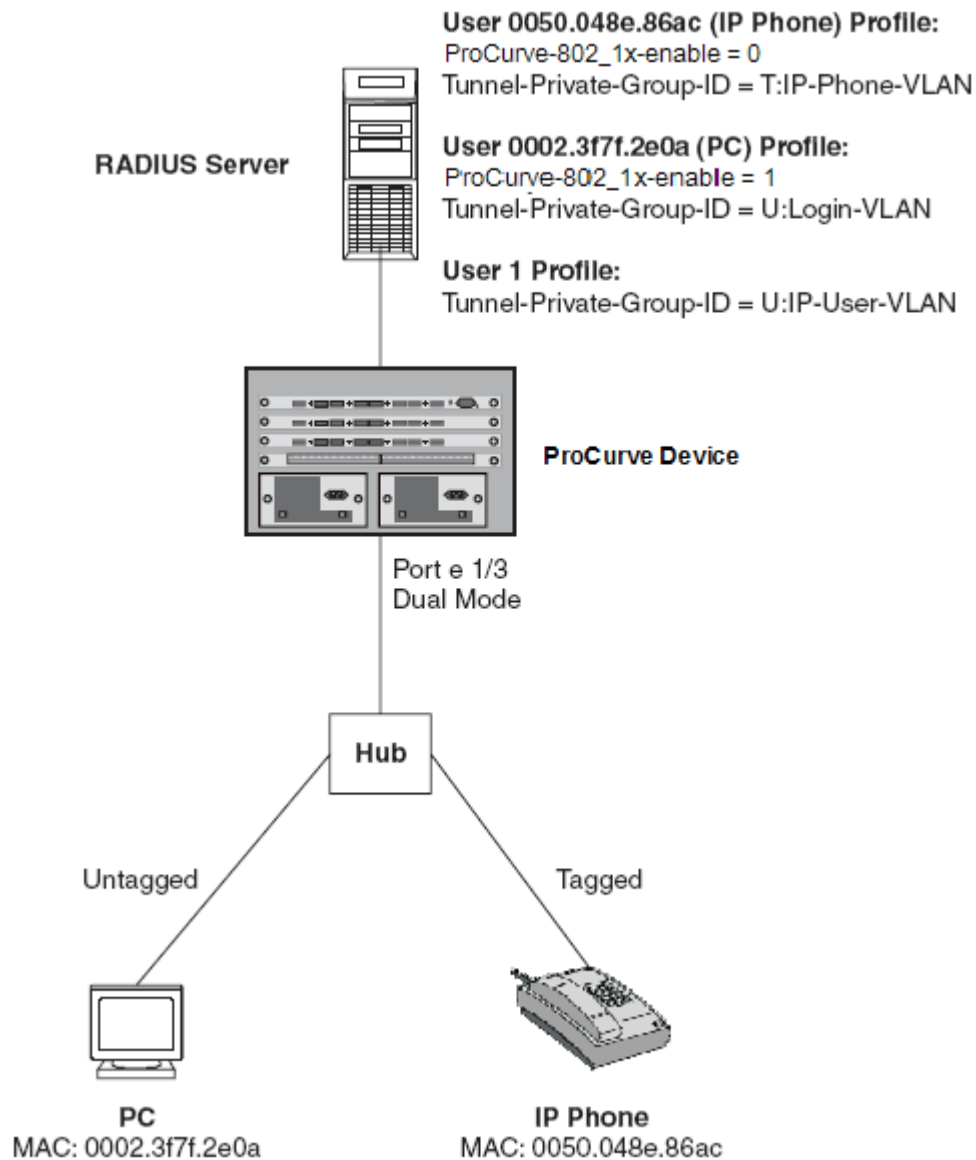


Figure 7 Sample configuration using multi-device port authentication and 802.1X authentication on the same port

When the devices attempt to connect to the network, they are first subject to multi-device port authentication.

When the IP phone's MAC address is authenticated, the Access-Accept message from the RADIUS server specifies that the IP phone's port be placed into the VLAN named "IP-Phone-VLAN", which is VLAN 7. The Foundry-802_1x-enable attribute is set to 0, meaning that 802.1X authentication is skipped for this MAC address. Port 1/3 is placed in VLAN 7 as a tagged port. No further authentication is performed.

When the PC's MAC address is authenticated, the Access-Accept message from the RADIUS server specifies that the PVID for the PC's port be changed to the VLAN named "Login-VLAN", which is VLAN 1024. The Foundry-802_1x-enable attribute is set to 1, meaning that 802.1X authentication is required for this MAC address. The PVID of the port 1/3 is temporarily changed to VLAN 1024, pending 802.1X authentication.

When User 1 attempts to connect to the network from the PC, he is subject to 802.1X authentication. If User 1 is successfully authenticated, the Access-Accept message from the RADIUS server specifies that the PVID for User 1's port be changed to the VLAN named "User-VLAN", which is VLAN 3. If 802.1X authentication for User 1 is unsuccessful, the PVID for port 1/3 is changed to that of the restricted VLAN, which is 1023, or untagged traffic from port e1 can be blocked in hardware.

Prior to authentication of the PC, the part of the running-config related to port 1/3 would be as follows:

```
interface ethernet 1/3
  dot1x port-control auto
  dual-mode
```

When the PC is authenticated using multi-device port authentication, the **dual-mode** statement for port 1/3 would be changed to reflect the dynamically assigned PVID of the "Login-VLAN", which is VLAN 1024:

```
interface ethernet 1/3
  dot1x port-control auto
  dual-mode 1024
```

After User 1 is authenticated using 802.1X authentication, the **dual-mode** statement for port 1/3 would be changed to reflect the dynamically assigned PVID of the "User-VLAN", which is VLAN 3:

```
interface ethernet 1/3
  dot1x port-control auto
  dual-mode 3
```

Example 2

The configuration in Figure 7 requires that you create a profile on the RADIUS server for each MAC address to which a device or user can connect to the network. In a large network, this can be difficult to implement and maintain.

As an alternative, you can create MAC address profiles only for those devices that do not support 802.1X authentication, such as IP phones and printers, and configure the ProCurve device to perform 802.1X authentication for the other devices that do not have MAC address profiles, such as user PCs. To do this, you configure the ProCurve device to perform 802.1X authentication when a device fails multi-device port authentication.

Figure 8 shows a configuration where multi-device port authentication is performed for an IP phone, and 802.1X authentication is performed for a user's PC. There is a profile on the RADIUS server for the IP phone's MAC address, but not for the PC's MAC address.

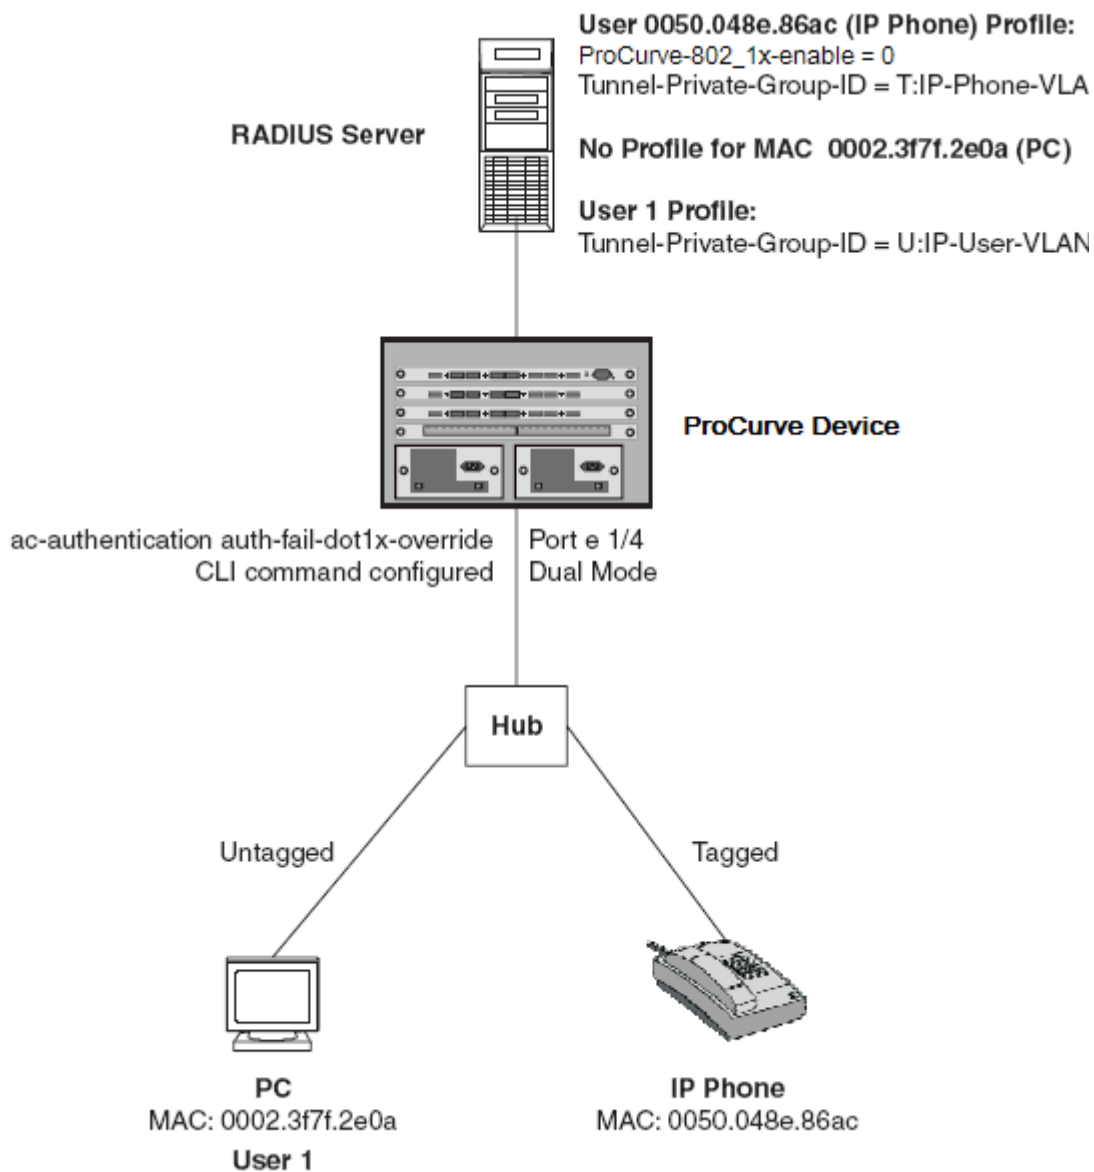


Figure 8 Sample configuration where 802.1X authentication is performed when a device fails multi-device port authentication

Multi-device port authentication is initially performed for both devices. The IP phone's MAC address has a profile on the RADIUS server. This profile indicates that 802.1X authentication should be skipped for this device, and that the device's port be placed into the VLAN named "IP-Phone-VLAN".

Since there is no profile for the PC's MAC address on the RADIUS server, multi-device port authentication for this MAC address fails. Ordinarily, this would mean that the PVID for the port would be changed to that of the restricted VLAN, or untagged traffic on the port would be blocked in hardware. However, the ProCurve device is configured to perform 802.1X authentication when a device fails multi-device port authentication, so when User 1 attempts to connect to the network from the PC, he is subject to 802.1X authentication. If User 1 is successfully authenticated, the PVID for port 1/4 is changed to the VLAN named "User-VLAN".

To configure the ProCurve device to perform 802.1X authentication when a device fails multi-device port authentication, enter the following command:

```
ProCurve 9300(config)# mac-authentication auth-fail-dot1x-override
```

Syntax: [no] mac-authentication auth-fail-dot1x-override

CPU ACL

ProCurve devices provide protection from denial of service (DoS) attacks for management protocols such as HTTP, SSH, Telnet, and SNMP. To do so, ACLs applied to ingress ports filter on these management and application protocols. In releases prior to 08.0.00, this feature is limited in that it provides protection from denial of service attacks for management protocols and applications only. It does not provide protection for all protocols.

Starting in release 08.0.00, you can configure the ProCurve device to provide protection from denial of service attacks for all protocols and applications received on the device. To do so, extended ACLs applied to ingress ports filter on the protocols specified in ACL clauses. For example, you can use extended ACLs to filter on protocols such as SNMP, TFTP, BGP, and ICMP.

The CPU ACL feature filters incoming packets in both hardware (EP only) and software. Packets filtered in hardware travel faster and use less resources in comparison to packets sent to the CPU for processing. If the device does not have available hardware resources (Content Addressable Memory (CAM)), the system will filter the packets in software (send the packets to the CPU for processing).

On a ProCurve Layer 3 switch, each interface can potentially be configured with multiple management IP addresses, in which case, CAM resources can quickly be consumed for all destination IP addresses. In this case, you can restrict management access to interfaces by disabling the management IP address through CAM (hardware denial).

The following sections provide instructions and examples for configuring CPU ACLs on a ProCurve Layer 3 Switch.

Configuring CPU ACL on a Routing Switch

NOTE: On a ProCurve Layer 3 switch, each interface can potentially be configured with multiple management IP addresses, thus, CAM resources can quickly be consumed for all destination IP addresses. In this case, you can restrict management access to interfaces by disabling the management IP address through CAM (hardware denial). See “Disabling an Interface’s Access to Management Functions” on page 76.

The following is an example of configuring an EP-based Layer 3 Switch to perform filtering for Telnet access.

```
ProCurve 9300(config)# vlan 3 by port
ProCurve 9300(config-vlan-3)# untagged ethe 3/1 to 3/5
ProCurve 9300(config-vlan-3)# exit

ProCurve 9300(config)# interface ve 3
ProCurve 9300(config-ve-3)# ip address 10.10.11.1 255.255.255.0
ProCurve 9300(config-ve-3)# exit

ProCurve 9300(config)# access-list 100 permit tcp host 10.10.11.254 any eq telnet
ProCurve 9300(config)# access-list 100 permit tcp host 192.168.2.254 any eq telnet
ProCurve 9300(config)# access-list 100 permit tcp host 192.168.12.254 any eq telnet
ProCurve 9300(config)# access-list 100 permit tcp host 192.64.22.254 any eq telnet
ProCurve 9300(config)# access-list 10 deny any
```

In this example, a Layer 3 VLAN is configured as a remote-access management VLAN and a router interface. The IP address specified for the router interface becomes the management IP address of the VLAN. Hardware filtering of Telnet traffic is performed for all the ports in the management VLAN. If CAM resources are exhausted, ACL clauses that are not programmed in CAM (hardware) will be filtered in software (sent to the CPU for processing).

When you make changes to the ACL configuration and/or make changes to the management VLAN, you must enter the following command after making the configuration changes:

```
ProCurve 9300(config)# remote-management rebind
```

Syntax: remote-management rebind

The **show cam I4** command displays the following information about the hardware filtering in this configuration:

```
ProCurve 9300(config)#show cam I4 3/1
Sl Index      Src IP_Addr  SPort      Dest IP_Addr  DPort Prot Age      Out Port
3 40960      192.64.22.254/32 Any        10.10.11.1/24 23 TCP dis Use L2/L3
3 40962      192.168.12.254/32 Any        10.10.11.1/24 23 TCP dis Use L2/L3
3 40964      192.168.2.254/32 Any        10.10.11.1/24 23 TCP dis Use L2/L3
3 40966      10.10.11.254/32 Any        10.10.11.1/24 23 TCP dis Use L2/L3
3 40968      Any          Any        10.10.11.1/24 23 TCP dis Discard
```

The IP address in extended ACL 100 is the source IP address of the filter entry, and the IP address of the router interface is always the destination IP address of the filter entry, regardless of the destination IP address specified in the ACL clause.

NOTE: Specifying a port number with the **show cam I4** command, for example, **show cam I4 3/2**, causes the device to display Layer 4 CAM information for all of the ports in the same port region (DMA) as port 3/2.

Example CPU ACL Configuration on a Layer 3 Switch

The following shows an example CPU ACL configuration.

```
ProCurve 9300(config)# ip receive access-list 100 vlan 3
ProCurve 9300(config)# access-list 100 deny icmp any any packet-too-big
ProCurve 9300(config)# access-list 100 permit icmp any any echo
ProCurve 9300(config)# access-list 100 permit udp 192.168.2/253/24 any eq tftp
ProCurve 9300(config)# access-list 100 permit udp 192.168.2.253/24 any eq snmp
ProCurve 9300(config)# access-list 100 permit udp 192.168.2.253/24 any eq snmp
ProCurve 9300(config)# access-list 100 permit tcp 192.168.2.253/24 any eq ftp
ProCurve 9300(config)# access-list 100 permit tcp 192.168.2.253/24 any eq http
ProCurve 9300(config)# access-list 100 permit tcp 192.168.2.253/24 any eq ssh
ProCurve 9300(config)# access-list 100 permit tcp 192.168.2.253/24 any eq telnet
ProCurve 9300(config)# access-list 100 permit ospf any any precedence internet
ProCurve 9300(config)# access-list 100 deny ip any any'
```

Disabling an Interface's Access to Management Functions

You can protect the CPU from remote access to management and application protocols. To enable this feature, disable access to the Management IP address through the device's Content Addressable Memory (CAM). The following shows an example configuration.

NOTE: This feature does not affect Layer 3 routing functions.

```
ProCurve 9300(config)# global-protocol-vlan
ProCurve 9300(config)# vlan 1 name DEFAULT-VLAN by port
ProCurve 9300(config-vlan-1)# exit

ProCurve 9300(config)# router ospf
ProCurve 9300(config-ospf-router)# area 0
ProCurve 9300(config-ospf-router)# exit

ProCurve 9300(config)# int e 3/10
ProCurve 9300(config-if-e1000-3/10)# ip address 10.10.10.1 255.255.255.0
```

```

ProCurve 9300(config-if-e1000-3/10)# ip ospf area 0
ProCurve 9300(config-if-e1000-3/10)# exit

ProCurve 9300(config)# int e 3/11
ProCurve 9300(config-if-e1000-3/11)# ip address 11.11.11.1 255.255.255.0
ProCurve 9300(config-if-e1000-3/11)# ip ospf area 0
ProCurve 9300(config-if-e1000-3/11)# management-ip-disable
ProCurve 9300(config-if-e1000-3/11)# exit

ProCurve 9300(config)# int e 3/12
ProCurve 9300(config-if-e1000-3/12)# ip address 12.12.12.1 255.255.255.0
ProCurve 9300(config-if-e1000-3/12)# ip ospf area 0
ProCurve 9300(config-if-e1000-3/12)# management-ip-disable
ProCurve 9300(config-if-e1000-3/12)# exit

ProCurve 9300(config)# int e 3/13
ProCurve 9300(config-if-e1000-3/13)# ip address 13.13.13.1 255.255.255.0
ProCurve 9300(config-if-e1000-3/13)# ip ospf area 0
ProCurve 9300(config-if-e1000-3/13)# management-ip-disable
ProCurve 9300(config-if-e1000-3/13)# exit

```

Syntax: [no] ip address <ip-addr> <ip-mask>

where <ip-addr> and <ip-mask> are the destination IP address and subnet mask.

Syntax: [no] management-ip-disable

Use the **no** form of the command to re-enable access to the Management IP address.

Viewing Information about Disabled Management IPs

Use the **show cam I4** command to display information about CAM entries for disabled Management IP addresses.

```
ProCurve 9300(config)#show cam I4 3/11
```

S1	Index	Src IP_Addr	SPort	Dest IP_Addr	DPort	Prot	Age	Out Port
3	40960	Any	Any	11.11.11.1/24	23	TCP	dis	Discard
3	40962	Any	Any	11.11.11.1/24	80	TCP	dis	Discard
3	40964	Any	Any	11.11.11.1/24	1812	TCP	dis	Discard
3	40966	Any	Any	11.11.11.1/24	49	TCP	dis	Discard
3	40968	Any	Any	11.11.11.1/24	22	TCP	dis	Discard
3	40970	Any	Any	12.12.12.1/24	23	TCP	dis	Discard
3	40972	Any	Any	12.12.12.1/24	80	TCP	dis	Discard
3	40974	Any	Any	12.12.12.1/24	1812	TCP	dis	Discard
3	40976	Any	Any	12.12.12.1/24	49	TCP	dis	Discard
3	40978	Any	Any	12.12.12.1/24	22	TCP	dis	Discard
3	43520	Any	Any	11.11.11.1/24	161	UDP	dis	Discard
3	43522	Any	Any	11.11.11.1/24	69	UDP	dis	Discard
3	43524	Any	Any	11.11.11.1/24	49	UDP	dis	Discard
3	43526	Any	Any	12.12.12.1/24	161	UDP	dis	Discard
3	43528	Any	Any	12.12.12.1/24	69	UDP	dis	Discard
3	43530	Any	Any	12.12.12.1/24	49	UDP	dis	Discard

See the *Command Line Interface Reference* for definitions of the fields shown in the display.

Where to Get More Information

These release notes provide information about features that are new since software version 07.8.01. For information about features in earlier software releases that are related to this release, see the documents listed in Table 1.

Software Fixes

The fixes for software bugs found on ProCurve 9300M Series routing switches in software releases 07.8.01 through 08.0.01k are described in the following table. Software fixes are listed by bug ID for the software release in which they were implemented.

Bug ID	Bug Description	Protocol/ Feature	Version Found	Version Fixed
	08.0.00			
28275	<p>On a 48-port RJ-45 10/100-TX module, the no rate-limit in command incorrectly removes the rate-limiting configuration from all ports that are configured to rate limit <i>inbound</i> traffic.</p> <p>Similarly, the no rate-limit out command incorrectly removes the rate-limiting configuration from all ports that are configured to rate limit <i>outbound</i> traffic. These commands should remove rate limiting only from the interface on which you enter the no rate-limit command.</p>	Rate-limiting		08.0.00
29086	On an 8-port Gigabit Ethernet Management module, an SNMP trap and Syslog message are not properly generated when the status of an LACP port changes from "up" to "block".	Link Aggregation		08.0.00
29530	The routing switch records several "VLAN table write error" messages in the system log.			08.0.00
29979	The web management interface does not allow you to configure named ACLs.	Web Management		08.0.00
31281	A routing switch reloads while running the "invalidate_buffer_length+136" function.	Other		08.0.00
31313	An error in a VRRP configuration on a virtual interface causes VRRP configuration errors in other virtual interfaces.	VRRP		08.0.00
31903 (also 33833)	On an 8-port Gigabit Management module, physically removing the transmit link on an MRP ring causes RHP packet loss. This causes other MRP rings in the same DMA to flap (fluctuate between "up" and "down").	MRP		08.0.00
32623	<p>On an 8-port Management module, a Network File System (NFS) configuration experiences slow performance after the queue depth is manually configured.</p> <p>Additionally, after configuring the queue depth then saving the configuration and reloading the device, the following error message is displayed on the console when the routing switch starts to adjust the queue depth:</p> <pre>Queue Depths value must be 0 to 15, line ignored!</pre>	Other		08.0.00

Bug ID	Bug Description	Protocol/ Feature	Version Found	Version Fixed
32776	If topology groups are enabled on a routing switch, the switch does not properly increment an SNMP get-Next or SNMP walk for the MIB object snVSrplfVlanId . Note that if you disable topology groups, this issue does not occur.	SNMP		08.0.00
32867	On an M4 24-port 10/100 module, you cannot access ports 13 to 24 using the front panel display on the web management interface.	Web Management Interface		08.0.00
33513	When a standard ACL is used to restrict Telnet access to the ProCurve device (using the telnet access-group <acl-number> command), the system log will show that the source port for both permit and deny ACL statements is always 0 (zero). Note that this problem occurs only with Telnet access. It does not occur when standard ACLs are used to restrict other access methods	Access Lists		08.0.00
33634	On an 8-port Gigabit Ethernet Management module, the running configuration displays no include ethernet x/x to y/y after you remove x/x and y/y trunk groups. For example, after you enter the no trunk ethernet 2/1 to 2/2 command, no include ethernet 2/1 to 2/2 is displayed in the running configuration file.	MSRP		08.0.00
34609	On a 2-port 10-Gigabit Ethernet module with Xenpak optics, a 10-Gigabit Ethernet link remains active and continues to forward traffic even though the other side of the link is disabled. In this particular configuration, Link Fault Signalling (LFS) is not enabled.	System		08.0.00
34823	Setting several system parameter limits to their maximum values (using the system-max command) can cause IP memory allocation errors if the memory in a routing switch cannot support all of the configured thresholds. When this occurs, the switch reloads and logs error messages to the console and the system log.	IP Stack		08.0.00
34843	A routing switch with the MAC Port Security feature is unable to learn a new secure MAC address. This problem occurs after a you enter the no secure-mac command at the Interface level to remove all learned secure MAC addresses. In this case, the switch does not immediately learn a new MAC address as a secure address unless the port is disabled for more than 60 seconds then re-enabled.	Port Security		08.0.00

Bug ID	Bug Description	Protocol/ Feature	Version Found	Version Fixed
34969	When you enter the no enable command at the Interface Configuration level, an interface or trunk group is not disabled correctly. Similarly, you enter the no disable command at the Interface Configuration level, an interface or trunk group is not enabled correctly.	CLI		08.0.00
35041	CLI and ProCurve documents state that the range of valid values for the default metric (default-metric <num> command) is a number from 1 through 15. However, the correct range of valid values are 1 through 65535; a routing switch will accept any number in this range. Starting in release 08.0.00, the CLI displays and accepts values 1 through 65535 with the default-metric command.	OSPF		08.0.00
35179	The web management interface (WMI) does not correctly display the currently configured Maximum Transmission Unit (MTU). In addition, the WMI does not support the ability to configure the default MTU nor to configure MTU at the port level; it supports the configuration of IP MTU only.	Web Management		08.0.00
35181	A routing switch incorrectly removes the padding from an IPX packet while forwarding the packet. Note that this problem occurs with EP modules. It does not occur with non-EP modules.	IPX Stack		08.0.00
35199	A routing switch reloads while forwarding sFlow packets with an invalid port number.	sFlow		08.0.00
35223	Some graphics on a routing switch do not load properly. For example, the thumbscrew graphics on ProCurve 9304M slots show "ACT" instead of the thumbscrew graphics. On a ProCurve 9304M with EP management module (J4885A), some of the port graphics do not display correctly.	Web Management Interface		08.0.00
35230	The global and per interface OSPF distribute-list filter does not correctly filter routes.	OSPF		08.0.00
35271	The no secure-MAC command causes a routing switch to convert a secure MAC address to a dynamic address in the MAC table. In this case, the switch should maintain the secure MAC address instead of converting it, even though an incoming violating packet on the MAC port was dynamically discovered.	Port Security		08.0.00
35278	Reverse Path Forwarding does not re-learn all dynamic and static routes after clearing the CAM entries for unicast RPF (clear ip rpf command) and clearing all IP routes from memory (clear ip route command).	Reverse Path Forward Check		08.0.00

Bug ID	Bug Description	Protocol/ Feature	Version Found	Version Fixed
35289	The first rate-limit entry in an ACL-based rate-limiting policy filters traffic correctly, but the second and subsequent rate-limit entries in the same ACL-based rate-limiting policy do not function correctly.	Rate Limiting		08.0.00
35290	A routing switch with Network Address Translation (NAT) incorrectly replies to a Telnet request instead of performing NAT for traffic between the outside NAT interface and inside NAT interface. This problem occurs if the next hop ARP address does not exist.	NAT		08.0.00
35292	LACP timer error messages occur during initialization of a fully populated routing switch that has over 100 aggregate links.	Link Aggregation		08.0.00
35328	If you enter the sa-filter in command to filter source-group pairs received in Source-Active (SA) messages from an MSDP neighbor, messages from all Source-Active caches are blocked (not only SA messages from peers that match the route-map) and are not displayed in show ip msdp sa command output. Also, when you enter the sa-filter originate command to filter source-group pairs in Source-Active messages in advertisements to an MSDP neighbor, the Source-Active cache is lost on other MSDP peers, not only on peers that match the route-map.	MSDP		08.0.00
35339	A routing switch incorrectly interprets a multicast SNAP packet as originating from a tunnel. As a result, the system displays the following error message before dropping the packet: Error! tunnel, but uptr_pkthdr=NULL	PIM Sparse		08.0.00
35512	A routing switch reloads when you enable an SSL server on the switch (using the web-management https command), and then enter the show run command at the CLI prompt	Other		08.0.00
35671	On an 8-port Gigabit Ethernet M4 Management module with mini-GBIC optics, an error in the software aging process causes the routing switch to reload.	Other		08.0.00
35799	Traffic flows incorrectly use Layer 4 flow-based CAM entries if the traffic flows do not match entries in an ACL-based rate limiting policy.	Rate Limiting		08.0.00
35906	Packet loss occurs intermittently when a transmit port is disconnected. This occurs in a configuration with Uni-directional Link Detection (UDLD).	UDLD		08.0.00

Bug ID	Bug Description	Protocol/ Feature	Version Found	Version Fixed
36040	A routing switch with Network Address Translation (NAT) cannot successfully complete ICMP echo requests (pings) from an inside host to another inside host. This problem occurs temporarily (lasts approximately six minutes) after the routing switch is reloaded. After six minutes, the pings are successful.	NAT		08.0.00
36108	In certain configurations, a routing switch may not properly flush OSPF summary LSAs (type 3 LSAs).	OSPF		08.0.00
36140	Entering the show tech command causes trunk ports to flap (fluctuate between “up” and “down”) when Uni-directional Link Detection (UDLD) is enabled.	System		08.0.00
36425	When an EP Management module is installed, an interface does not respond to a ping if you configure ACL-based rate limiting on the interface, then save the configuration and reload the routing switch. If you remove the rate limiting configuration, the interface will respond to a ping.	Rate limiting		08.0.00
36440	A routing switch incorrectly forwards IEEE 802.1D MAC bridge group addresses 01-80-C2-00-00-01 to 01-80-C2-00-00-0F.	N/A		08.0.00
36473	Adding or deleting a port on a VLAN that has one or more virtual interfaces may cause network connectivity issues. If you disable then re-enable the virtual interfaces, network connectivity is re-established.	VRRP		08.0.00
36769	When a 4-port Gigabit Ethernet EP Management module is installed, a routing switch reloads and fluctuates between the active management module and the standby management module while removing an ACL from an interface. This persists four or five times before the ACL is successfully removed from the interface.	Access Lists		08.0.00
36786	A GI Buffer Error message is displayed on the console when an 802.1X authentication session is running. 802.1X authentication creates a buffer overrun when it sends a packet size larger than the maximum size supported in the real buffers.	Other		08.0.00
37168	Port mirroring of traffic from one port to another is supported only in one direction if both ports are fiber ports and both are configured as monitor ports. Note that this problem does not occur on 10/100 ports.	Mirroring		08.0.00

Bug ID	Bug Description	Protocol/ Feature	Version Found	Version Fixed
37170	A routing switch reloads after a virtual interface is configured to disable access to the Management IP address (CLI command management-ip-disable) and ToS-based QoS (CLI command qos-tos) is enabled on the interface.	QoS		08.0.00
37171	When an 8-port Gigabit Ethernet EP Management module is installed, a routing switch reloads after taking more than 30 seconds to complete the CLI command show ip net-aggregate .	CLI		08.0.00
37244	A routing switch incorrectly flushes a network Link State Advertisement (LSA, type 2) after port flaps cause the router ID to change.	OSPF		08.0.00
37387	A routing switch reloads after a PRAM write error for a module that does not exist in the chassis.	Other		08.0.00
37424	When an EP Management module is installed, a routing switch forwards broadcast traffic to an invalid port, which in turn causes a system reload.	Other		08.0.00
37557	A memory leak may occur if the MIB object snFdbTableStationFlush is set to an invalid value.	SNMP Management		08.0.00
37644	The show version command displayed a corrupted serial number for a secondary management module.	CLI		08.0.00
37653	When you enter the no include-port ethernet command in software releases 07.6.06f and 07.6.06g, the running configuration is not displayed in show running-config command output.	VSRP	07.6.06f	08.0.00
37797	Although rare in occurrence, a routing switch may encounter a stack overflow after loading the configuration file.	System		08.0.00
37853	An invalid CAM index entry causes a system reload.	Other		08.0.00
37897	When a T-Flow Management module is installed, policy-based routing does not work after a default route is applied.	Policy Based Routing		08.0.00
37900	After you enable RIP on an interface, the web management interface does not detect a configuration change when you disable RIP. For example, if you disable RIP through the CLI, the show run command output shows that RIP is no longer enabled. However, the web management interface incorrectly displays RIP as "enabled".	Web Management		08.0.00

Bug ID	Bug Description	Protocol/ Feature	Version Found	Version Fixed
37911	<p>A routing switch incorrectly transposes the parameters for the radius-server <hostname> command. For example, if you enter the following CLI command:</p> <pre>radius-server host www auth-port 1812 acct-port 1813 default key test</pre> <p>The system transposes the command parameters as follows:</p> <pre>radius-server host 61.200.194.166 auth-port 1812 acct-port 1813 default key 0 www</pre> <p>This can be viewed in the output of the show-running config command.</p>	RADIUS		08.0.00
37921	<p>If you specify a large, invalid value with the CLI command topology-group, the returned error message incorrectly displays a negative value. For example:</p> <pre>ProCurve(config)# topology-group 4294967295 Error - -1 not between 1 and 256</pre> <pre>ProCurve(config)# topology-group 2147483649 Error - -2147483647 not between 1 and 256</pre>	Topology Groups		08.0.00
38062	A routing switch reloads after a 48-port Telco forwarding module is physically inserted in the chassis.	Other		08.0.00
38066	In a configuration with two firewalls and two ports with static MAC addresses, the device fails to forward packets from one port to the other.	Other		08.0.00
38070	An ACL incorrectly drops ICMP packet replies from a host. This occurs with certain bit settings in the ICMP packet header.	Other		08.0.00
38135	The system incorrectly resets a Gigabit Ethernet port on a non-EP module.	Other		08.0.00
38138	A routing switch does not properly resolve an ARP for a host MAC address in a local network. This causes network connectivity issues with VRRP-E.	Other		08.0.00
38198	<p>The following error message appears on the console, but is not written to the Syslog:</p> <pre>WARNING: Can't allocate routing table due to CAM 5 hardware error.</pre> <p>By default, all messages should be sent to the system log.</p>	Syslog		08.0.00

Bug ID	Bug Description	Protocol/ Feature	Version Found	Version Fixed
38270	A routing switch recognizes only 15 trunk groups, even though there are more than 15 trunk groups configured on the switch. The output of the show trunk command lists the correct number of <i>configured trunks</i> (for example, 18 trunks), but lists only 15 <i>operational trunks</i> . In addition, trunk groups greater than 15 do not inherit the sFlow configuration, even though sFlow is enabled on all ports.	Trunking		08.0.00
38315	Policy Based Routing (PBR) stops working when a routing switch inadvertently removes a Layer 4 CAM entry after a port fluctuates between “up” and “down”.	IP Policy		08.0.00
38332	The web management interface accepts an invalid trunk configuration instead of rejecting it. In addition, the routing switch continually reboots if the invalid trunk configuration is saved and the switch is reloaded. To fix this problem, you must reboot the routing switch with the default configuration.	Web Management		08.0.00
38674	An invalid calculation in the dynamic port data structure causes a system reload.	Other		08.0.00
39537	OSPF does not properly converge in an OSPF Not-So-Stubby Area (NSSA) after a link fluctuates between “up” and “down” states.	OSPF		08.0.00
39562	In a configuration with multiple ACL-based rate limiting policies, the most recently configured ACL-based rate limiting policy does not filter traffic. This problem occurs if you remove an ACL-based rate limiting policy then re-apply it.	Rate Limiting		08.0.00
39565	A routing switch applies the first 52 ACL-based rate limiting policies, even though there are 102 policies configured on the switch.	Rate Limiting		08.0.00
39584	A routing switch may drop some ICMP fragments if an inbound ACL permit statement specifies the ICMP protocol type.	Access Lists		08.0.00
39637	A routing switch reloads during multicast hardware verification when the <code>mcast_tx_entry_addr_verify</code> instruction is executed.	Other		08.0.00
38755	A routing switch incorrectly prefers RIP routes over OSPF intra-area routes, even though the switch is configured to redistribute RIP routes into OSPF. In addition, OSPF intra-area routes incorrectly inherit a non-zero Type 2 Cost (<code>Type2_Cost</code>).	OSPF		08.0.00

Bug ID	Bug Description	Protocol/ Feature	Version Found	Version Fixed
39782	In software release 07.8.00, the CLI command hardware-drop enables a routing switch to drop unwanted PIM Dense or PIM Sparse multicast packets in hardware, if the PIM session has no output interface. However, if PIM Dense or PIM Sparse packets are received through a tunnel, this feature may cause a system reload.	Other	07.8.00	08.0.00
39923	Multicast Source Discovery Protocol (MSDP) peers incorrectly advertise source addresses (SAs) with an incorrect PIM Sparse Rendezvous Point (RP) as originator.	MSDP		08.0.00
40018	A routing switch ignores the logging facility command if there is additional white space at the end of the command. In addition, if you use the Tab key to complete the command line or to search for additional options, the switch automatically appends extraneous white space to the end of each line. This also causes the switch to ignore the logging facility command.	CLI		08.0.00
40053	A routing switch reloads when the code reached the end of the ARP table, but instead of breaking out, the code continued into the access write-protect area.	Other		08.0.00
40093	A routing switch reloads during a routine AppleTalk procedure.	AppleTalk		08.0.00
40149	When port monitoring is enabled on an interface, the interface is not able to receive incoming traffic from another routing switch. If port monitoring is disabled on an interface, the interface is again able to receive incoming traffic from the other routing switch.	Monitor		08.0.00
40150	The output of the show run command displays an incorrect VLAN configuration after you add a VLAN to a VLAN group. For example, the command add-vlan 181 to 200 places the entry add-vlan 16981 in the running configuration.	VLAN by Port		08.0.00
40663	A Gigabit Ethernet uplink incorrectly forwards packets with a VLAN ID of zero (0), causing the port to lock up.	Uplink		08.0.00
40717	A system reload occurs due to a data access exception on a routing switch running software release 07.6.04a,.	Other		08.0.00
40718	A system reload occurs during a Telnet session due to program exception errors.	Telnet		08.0.00
40854	You cannot use SNMP to add a port from one VLAN to another VLAN if one of the VLANs was discovered by GVRP.	SNMP		08.0.00

Bug ID	Bug Description	Protocol/ Feature	Version Found	Version Fixed
40857	<p>You cannot add VLANs to a VLAN group if the VLANs are added in a certain order. For example, you cannot add the VLANs in the following order:</p> <pre>ProCurve 9300(config-vlan-group-1)#add-vlan 100 to 300 ProCurve 9300(config-vlan-group-1)#add-vlan 700 to 900</pre> <p>However, you can add VLANs in the following order:</p> <pre>ProCurve 9300(config-vlan-group-1)#add-vlan 700 to 900 ProCurve 9300(config-vlan-group-1)#add-vlan 100 to 300</pre>	Topology Groups		08.0.00
40948	Warning messages are displayed when a 16-port Gigabit Ethernet EP module is inserted in a routing switch running software release 07.6.06x.	Hardware		08.0.00
41001	A system reload occurs when the routing switch uses source port zero (0) for NAT when a slot in the switch is empty.	NAT		08.0.00
41003	If SSH with RSA authentication is enabled, an AAA authentication login at a privilege level is not supported.	AAA		08.0.00
41074	A routing switch reloads when Reverse Path Forwarding is enabled (CLI command ip verify unicast reverse-path).	Reverse Path Forward Check		08.0.00
41271	A routing switch does not program an implicit ACL deny statement in the Layer 4 CAM. As a result, the switch allows traffic instead of denying it.	Access Lists		08.0.00
41411	The show interface brief command output incorrectly shows that a port is running at 100 Mbps in full duplex when the port is running at 100 Mbps at half duplex. This occurs after a port on an EP 48-port 10/100 TX module auto-negotiates to 100 Mbps operating at half duplex (100-half), and then is manually configured to 100-full using the CLI command speed 100-full .	Auto Negotiation		08.0.00
41426	On an EP 48-port RJ-45 Forwarding module, a port on the forwarding module locks up when operating in half duplex or full duplex mode.	Auto Negotiation		08.0.00
41572	A routing switch fails to re-distribute the load-sharing on the default route after a Shortest Path First (SPF) calculation/route table update. This problem occurs when Default Route Aggregation is configured on the switch (CLI command ip dr-aggregate).	ECMP		08.0.00
41582	A routing switch reloads after trying to delete a nonexistent DMA CAM entry.	Other		08.0.00

Bug ID	Bug Description	Protocol/ Feature	Version Found	Version Fixed
41588	A routing switch stops sending PIM registration packets for source group (S,G) pairs. As a result, source groups pairs are not added to the Rendezvous Point's (RP's) multicast forwarding cache.	PIM Sparse		08.0.00
41614	A routing switch fails to record a trunk port status change in the running configuration file. This occurs if the change was made through SNMP.	SNMP Management		08.0.00
41640	When you configure the software aging period for blocked MAC addresses (CLI command mac-authentication max-age), the CLI does not return the range of possible max-age values after you enter "?" and press Return.	MAC Security		08.0.00
41898	An invalid port data structure in Layer 2 NetFlow causes a system reload.	Other		08.0.00
41989	An OSPF Autonomous System Boundary Router (ASBR) with external route summarization has delays when re-converging with an alternate path. This problem occurs when an area range is configured on a routing switch.	OSPF		08.0.00
42006	A routing switch may not properly clear CAM and cache entries after a new equal-cost path is added. This occurs in a configuration with Equal-Cost Multi-Path (ECMP) and more than two equal-cost paths.	ECMP		08.0.00
42047	An interface configured with advanced QoS features stops responding to ICMP echo requests (pings). This occurs after an incoming ACL is configured on another interface.	QoS		08.0.00
42051	A system reload occurs when you configure a routing switch.	OSPF		08.0.00
42115	A routing switch fails to dynamically configure a Voice over IP (VoIP) phone when it is connected to the device. This occurs even though a voice VLAN ID is configured on the port to which the VoIP phone is connected.	Other		08.0.00
42129	A routing switch sends debugging output to the console, even though the configured debugging output destination is Telnet only.	CLI		08.0.00
42150	A port's ACL-based rate-limiting policy incorrectly prevents the port from establishing TCP connections with other ports that do not have the same rate-limiting policy. This problem occurs after you enter the clear ARP command.	Rate Limiting		08.0.00
42156	Adjacency fails when OSPF sends a large database packet with MD5 authentication to a Netscreen router.	OSPF		08.0.00

Bug ID	Bug Description	Protocol/ Feature	Version Found	Version Fixed
43150	When you enter the no port-name command at the Interface level, the command may not remove the port name and may modify the existing port name.	CLI		08.0.00
43161	The output of the show authenticated-mac-address command displays an incorrect MAC authentication age for a MAC address. This field applies to unauthorized entries only and, in this case, should display a value of "N/A".	MAC Authorization		08.0.00
43216	The no rp-candidate add CLI command corrupts the PIM router configuration.	PIM Sparse		08.0.00
43300	An internal packet processing function causes a routing switch to reload.	EP Management Module		08.0.00
43385	On a 48-port 10/100-TX module, a multicast listener stops receiving multicast packets if IGMP timers are reconfigured to a value of more than 255.	IGMP		08.0.00
43424	A routing switch fails to update the OSPF NSSA default route with an alternate path after the primary path goes down. Re-initializing the OSPF session temporarily resolves this problem.	OSPF		08.0.00
43482	A routing switch does not update the MAC and ARP tables when an incoming hardware-based ACL is configured on a virtual interface. To temporarily fix this problem, enter the clear mac and clear arp commands and then remove and re-add the ACL on the virtual interface.	Access Lists		08.0.00
43651	A routing switch does not properly forward multicast traffic within a multi-slot trunk group. Instead, it sends these packets to the CPU where they are eventually dropped. To temporarily resolve this problem, remove the multi-slot configuration from the slot on which the multicast packets originate, or disable route-only on the trunk group.	Trunking		08.0.00
43653	The SNMP View-based Access Control Model (VACM) tables display the read-write community string configured on the device. The read-write community string should not be displayed.	SNMP Management		08.0.00

Bug ID	Bug Description	Protocol/ Feature	Version Found	Version Fixed
43740	The Syslog shows invalid Transmission Control Block (TCB) access error messages.	TCP Stack		08.0.00
43910	IGMP V3 fast leave waited three seconds to stop traffic when a client changed state from "excluding none" to "including none". In normal operation, the routing switch stops traffic immediately after you enter the ip igmp tracking command on a virtual interface to enable tracking and the fast leave feature.	IGMP		08.0.00
44041, 44042	On an 8-port mini-GBIC module, a system reload occurs when you enter the show run command.	SSH		08.0.00
44138	CPU utilization increases for about 8 seconds when ports are added to or removed from a VLAN that has outbound ACLs configured.	ACL		08.0.00
44254	SNMP reports a "0" value when SNMP is used to query the VRRP dead-interval. In normal operation, the default or configured value is reported.	VRRP		08.0.00
44293	RADIUS notification messages are not sent, resulting in the failure of an authentication session.	Authentication		08.0.00
44313	A routing switch sends a failure message at the start of authentication, resulting in failed authentication attempts.	Authentication		08.0.00
44439	MTU field must be set to "0" in order to be compatible with RFC 1583.	OSPF		08.0.00
44515	Packets are dropped because the MAC table was not flushed after an 802.1W topology change.	RSTP		08.0.00

Bug ID	Bug Description	Protocol/ Feature	Version Found	Version Fixed
	08.0.01b			
54927	When RSTP is configured, the transition from the designated port to the root port requires two seconds after a link failure occurs on the root port.	RSTP		08.0.01b
55534	<p>The routing switch uses the lowest next hop of multi-path routes in the reverse path forwarding (RPF) check.</p> <p>Note: After you activate PIM multicast by entering the router pim command, you must enable the Highest IP RPF feature (highest-ip-rpf command) for the shortest path back to the source to be selected. The selection is based on which Reverse Path Forwarding (RPF) neighbor in the IP routing table has the highest IP address.</p> <p>Enabling the Highest IP RPF feature avoids route confusion between network devices in a complicated multi-path topology. This feature is enabled in the default configuration in release 08.0.01b. If your network has ProCurve routing switches running earlier software releases, it is recommended that you enable the Highest IP RPF feature on older software releases.</p> <p>You can disable the Highest IP RPF feature by entering the no highest-ip-rpf command at PIM router configuration level.</p> <p>Note: The routing switch configuration does not display "highest-ip-rpf" when this feature is enabled by default. The configuration only displays "no highest-ip-rpf" when the Highest IP RPF feature is disabled.</p>	Multicast		08.0.01b
56466	When you enter the show span detail command to display detailed Multiple Spanning Tree (MSTP) information for a nonexistent individual port, a system reload occurs.	MSTP		08.0.01b

Bug ID	Bug Description	Protocol/ Feature	Version Found	Version Fixed
58712	<p>When RSTP is configured and the port state on the inactive path changes from “discarding” to “forwarding”, a Topology Change (TC) message is sent out. The processing of a TC message results in a MAC flush that may result in transient packet loss.</p> <p>To disable the sending of Topology Change messages, enter the following command for each VLAN on the inactive path:</p> <p>span 802-1w no-tc-on-inactive-path-converge</p> <p>When you enter this command in a point-to-point topology, the sending of the TC message is disabled.</p> <p>When you enter this command on a designated port in an RSTP configuration with point-to-point links, in which the Proposal/Agreement sequence is not used, the sending of a TC message is disabled.</p> <p>In all other configurations, the sending of a TC message is not disabled.</p>	RSTP		08.0.01b
	08.0.01c			
52810	If an ACL with DSCP mapping is applied on a virtual routing interface (VE), the ACL incorrectly matches the VLAN., causing OSPF packets sent on the VLAN to be discarded.	OSPF		08.0.01c
56905	When you enter the mstp inst 0 ethernet 3/2 path-cet 22000 command, a system reload occurs if interface 3/2 is not present.	Spanning Tree		08.0.01c
58852	<p>In earlier software releases, a BGP community attribute with an extended length greater than 256 bytes was not processed properly.</p> <p>A BGP community attribute with an extended length greater than 256 bytes is now supported.</p>	BGP		08.0.01c
59604	When CPU Protection is enabled on a VLAN that has one or more trunks configured and a member port is on the same DMA as the trunk, the packets received on the DMA are dropped.	CPU Protection		08.0.01c
	08.0.01d			
56816	When an ACL with qos-tos mapping is configured on a virtual routing (VE) interface, the VLAN is not updated. As a result, some packets are incorrectly associated with the wrong VLAN and are discarded.	ACL		08.0.01d

Bug ID	Bug Description	Protocol/ Feature	Version Found	Version Fixed
60117	When VSRP is enabled on a VLAN, its port state is not correctly reflected as "Blocked" in the Layer 2 VLAN. In such cases, packets received on a port may cause a loop.	VSRP		08.0.01d
	08.0.01e			
44137	CAM entries are not created for "isolated" VLANs with untagged ports on a routing switch.	Private VLAN		08.0.01e
56078	Although the range of valid values for the VRRP Slow Start feature are from 1 to 255, a routing switch accepts a value outside this range.	VRRPE		08.0.01e
59468	Reconfigured OSPF settings that influence interface configurations are not displayed in show ip ospf config command output. You must reboot the routing switch before the OSPF changes are displayed in the output.	OSPF		08.0.01e
59519	When OSPF is configured on a fixed full-duplex (10-Full or 100-Full) interface, the external route summarization fails when the configuration is loaded during system reload.	OSPF		08.0.01e
60577	If Uni-Directional Link Detection (UDLD) is configured, UDLD packets may get lost, resulting in link a "down" state.	UDLD		08.0.01e
61237	If SSH is configured, memory corruption may occur when an incorrect SSH key length is received.	SSH		08.0.01e
	08.0.01f			
43428	If you access a routing switch using HTTP, a system reload may occur when an internal request to allocate a buffer of 8K or greater is made.	Web Management Interface		08.0.01f
54219	The default SSL HTTPS certificate for the Web Management interface expires on 12/31/05. HTTPS sessions will start to fail on 1/1/06.	Web Management Interface		08.0.01f
57291	In a routing switch running IPv6 with OSPF, the OSPF timer does not shut down correctly, resulting in a system reload.	OSPF (IPv6)		08.0.01f
59070	If you apply policy-based routing (PBR) to a virtual interface, the traffic is not forwarded correctly to the appropriate next hop.	PBR		08.0.01f
60396	If a static MAC address is configured on multiple ports, the list of MAC addresses displayed in show mac-address ethernet <portnum> command output contains incorrect information.	System		08.0.01f
61175	MAC addresses are not learned when a tagged port is reconfigured as an untagged port.	System		08.0.01f

Bug ID	Bug Description	Protocol/ Feature	Version Found	Version Fixed
61456	If the Multicast Source Discovery Protocol (MSDP) is configured, the route map used to filter the multicast source active (SA) cache does not filter source IP addresses correctly.	MSDP		08.0.01f
62367	If the Virtual Switch Redundancy Protocol (VSRP) is configured, the "link up" notification causes port 1/1 to become non-operational. As a result, no outbound traffic is sent on the port.	VSRP		08.0.01f
62887	If OSPF is configured and you modify the OSPF cost configured on a 1-Gigabit interface, the new cost value is not processed and the correct cost is not applied.	OSPF		08.0.01f
63310	If Protocol Independent Multicast sparse mode (PIM-SM) snooping is configured, the IP address and other configuration information provided by a DHCP server is not forwarded to a client.	IP Stack		08.0.01f
08.0.01g				
45575	If MAC authentication is configured, the NAS_PORT attribute is not sent to the RADIUS server. Instead, the routing switch sends the NAS_PORT_TYPE attribute.	AAA		08.0.01g
48697	If CPU protection is configured for flooding multicast traffic, multicast packets are sometimes discarded.	Port Security		08.0.01g
48698	To turn off the automatic tunneling of Cisco Systems' Per VLAN Spanning Tree (PVST) BPDUs across a Layer 2 domain, the pvst-auto-disable command has been implemented. Enter this command at the interface level; for example: ProCurve(config)# interface ethernet 1/1 ProCurve(config-if-e10000-1/1)#pvst-auto-disable	STP		08.0.01g
52419	Secure Shell (SSH) sessions are timed out even when an SSH session is active and in use.	SSH		08.0.01g
54219	The default Secure Sockets Layer (SSL) HTTPS certificate for Web management expires on 1-1-06. HTTPS sessions using the new certificate no longer fail.	HTTPS		08.0.01g
54250	If the Spanning Tree Protocol (STP) is configured, the debug span all command is not parsed correctly.	STP		08.0.01g
54251	If the Rapid Spanning Tree Protocol (RSTP) is configured, the debug 802.1w all <vlan> command does not cause debug messages to be sent to the console.	RSTP		08.0.01g

Bug ID	Bug Description	Protocol/ Feature	Version Found	Version Fixed
56798	A port security violation on one VLAN may affect the traffic on other VLANs.	Port Security		08.0.01g
57245	The aging out of Content Addressable Memory (CAM) entries may result in the incorrect deletion of an entry in the CAM table, resulting in a system reload.	Port Security		08.0.01g
57315	A secure MAC address supported on different VLANs does not age out with the same time interval.	Port Security		08.0.01g
59070	If policy-based routing (PBR) is configured with two or more instances of a PBR route map, the policy is applied only to the first instance.	PBR		08.0.01g
59793	If ACLs are configured and an ACL filters an invalid source port in an internal packet header, a system reload may occur.	System		08.0.01g
59925	The information associated with a CAM entry that is stored in the CAM table may cause a system reload.	System		08.0.01g
60227	If you configure interface speed on a port that does not exist, a system reload may result.	System		08.0.01g
60401	<p>In earlier software releases, when you configure a password for a username using all numeric characters, you must re-enter the service password-encryption command after you add each password.</p> <p>The create-password keyword is now supported in the following username commands to create a password without having to enter the service password-encryption command:</p> <ul style="list-style-type: none"> username <username> create-password <password-string> <p>Or:</p> <ul style="list-style-type: none"> username <username> privilege <number> create-password <password-string> <p>Note that you can generate the same password by using the existing username command syntax:</p> <ul style="list-style-type: none"> username <username> password <number> <encryptedpassword-string> <p>Or</p> <ul style="list-style-type: none"> username <username> privilege <number> password <number> <encrypted-password-string> 	AAA		08.0.01g

Bug ID	Bug Description	Protocol/ Feature	Version Found	Version Fixed
61336	If CPU protection is configured to flood unicast packets, source MAC addresses are not learned, causing traffic to be sent to the CPU and resulting in packet loss for traffic with higher rates.	System		08.0.01g
61501	If the Spanning Tree feature is configured, a loop in the default VLAN occurs when a port that belongs to the default VLAN and has the following configuration settings is added to another VLAN: <ul style="list-style-type: none"> The port is configured to operate as a dual-mode port. The port is configured with an STP state set to "always forwarding". 	STP		08.0.01g
62862	If RIP routes that are redistributed to OSPF are configured, when an interface that learns RIP routes is disabled momentarily, the redistributed RIP routes are correctly removed but are not restored when the interface comes back up.	OSPF		08.0.01g
62887	If you reconfigure the OSPF link cost on a 1-Gigabit interface, the interface cost configuration is not updated correctly.	System		08.0.01g
62893	When a port security violation brings down a trunk's primary port, the show port security mac command output incorrectly displays the secondary port state.	Port Security		08.0.01g
62894	When a port security violation causes a trunk to be disabled, you must enter the enable command at the trunk configuration level to recover the trunk. When the trunk is disabled, its "disable" status is not displayed in trunk group configuration information.	Port Security		08.0.01g
63329	If the Port Security feature is configured, when a MAC address violation occurs on the secondary member of a port, the secondary port is disabled even if you do not enter the trunk-config-ind command at the trunk group configuration level.	Port Security		08.0.01g
63629	When OSPF is configured, the last OSPF route that is used is listed twice in show ip ospf routes command output.	OSPF		08.0.01g
63837	If VSRP is configured and a new port is added to the VSRP domain, the VSRP CAM entry is not programmed, causing packets received on the port to be discarded.	VSRP		08.0.01g
63921	If UDLD is configured, an attempt to access an invalid port number may cause a system reload.	UDLD		08.0.01g

Bug ID	Bug Description	Protocol/ Feature	Version Found	Version Fixed
64605	If you enter the tag-type command and later delete the configured tag-type value, the exact hexadecimal value is not checked. You can enter any hexadecimal number to delete the configured value.	System		08.0.01g
64621	If SSHv2 is configured with AAA authorization and accounting as the default, you must press the Enter key twice before the password prompt is displayed for an SSH session.	SSH		08.0.01g
64666	If you access the MIB hierarchy using SNMP walk, an SNMP getnext request does not increase the object identifier (OID) and generates an error.	SNMP		08.0.01g
64871	If policy-based routing is configured and more than one interface is configured with a PBR policy, the ip rebind-acl all command does not rebind ACL clauses for PBR.	PBR		08.0.01g
64869	The system information displayed in show version command output using SNMP does not display the backplane type in the system description, although this information is displayed when you enter the show version command from the command-line interface.	System		08.0.01g
65585	EP 16-Port Gigabit Ethernet modules do not auto-negotiate flow control capability correctly.	System		08.0.01g
	08.0.01h			
48833	If SSH is configured, the debug ip ssh level command, which is used to troubleshoot an SSH connection to the routing switch, is not supported.	SSH		08.0.01h
63429	When you download a group of ACLs to the running configuration using the copy tftp running-config command, a system reload may occur.	ACL		08.0.01h
65957	If a large number of virtual routing interfaces are configured with the Policy Based Routing (PBR) feature, access to a console session following a system reload may take a long time.	PBR		08.0.01h
66238	When a SuperSpan boundary is configured on a 2-Port 10 Gigabit Ethernet module with Xenpak optics, packets are sometimes dropped.	SuperSpan		08.0.01h
	08.0.01i			
55771	If OSPF is configured, the summary route continues to be advertised although the source IP address used in the summary-route advertisement is removed.	OSPF		08.0.01i

Bug ID	Bug Description	Protocol/ Feature	Version Found	Version Fixed
66207	On a a 10 Gigabit Ethernet module, if you enter the debug hw command to troubleshoot hardware backplane polling, a system reload may occur.	System		08.0.01i
66840	Starting in release 08.0.01i, the phy x10g xenpak vendor command is supported on a 10 Gigabit Ethernet module with Xenpack optics to display vendor-specific information, such as name, serial number and manufacturing date; for example: ProCurve# phy 4/1 x10g 18 xenpak vendor XENPAK:PMA id 1=00000182 PMA id 2=00003c00 XENPAK std rev. support: 30 Vendor Name: Optillion AB Part Number: TOP 3010- SC Revision: B Serial Number: KK03260028 Year-Month-Day: 20030602	System		08.0.01i
67005	On a a 10 Gigabit Ethernet module, the same packet may be forwarded twice: the first time by the hardware, the second time by the CPU.	System		08.0.01i
67109	If RIPv2 and the redistribution of static routes are configured, disabling an interface on which a next hop is configured causes RIP to incorrectly advertise a static route, without changing the metric value to 16.	RIP		08.0.01i
67841	If a trunk, UDLD, and RSTP are configured in a loop topology, a port that is in a discarding state sometimes causes a UDLD link to flap.	UDLD		08.0.01i
08.0.01j				
60915	When a BGP peer group is configured without a remote autonomous system (AS) on a routing switch, the configuration is not saved after a system reload.	BGP		08.0.01j
63830	If you disable the VRRP master port on a routing switch configured with VRRP, the switch sends out a gratuitous ARP packet with the physical MAC address.	VRRP		08.0.01j
66956	On the tagged VLAN ports of a 10 Gigabit Ethernet module, the jumbo MTU setting does not take into account the increase in packet length caused by VLAN tagging.	Stack		08.0.01j
67848	The show ip interface <slot/port> command does not display the correct IP MTU value configured on the routing switch.	Stack		08.0.01j

Bug ID	Bug Description	Protocol/ Feature	Version Found	Version Fixed
67854	If you modify the MTU value that was configured with the ip mtu command, the new MTU value is not updated.	Stack		08.0.01j
68140	Changes in the state of an OSPF neighbor are not logged if traps are enabled and the affected interface is not the designated router (DR).	OSPF		08.0.01j
68191	The display buffers used for CLI command output are sometimes lost when you enter the show tech command.	System		08.0.01j
68363	An MSDP peer may not be re-established if the interface used to reach the peer is disabled and re-enabled, and if the routing switch is using an indirect route to reach the MSDP peer.	MSDP		08.0.01j
69140	A system reload may occur if you enter the show span when a single spanning tree with a large number (several hundred) of VLANs is configured on a routing switch.	Spanning Tree		08.0.01j
	08.0.01k			
49848	If sFlow is configured, exporting packet counters and traffic statistics to an external collector may result in a system reload.	System		08.0.01k
55268	When you use the web management interface to configure a routing switch, a system reload may occur.	Web Management		08.0.01k
59925	A loop in the Content Addressable Memory (CAM) list may be caused when you delete CAM entries, resulting in a system reload.	System		08.0.01k
61205	If Network Address Translation (NAT) and ACLs are configured on a routing switch, an ACL configured to filter outbound traffic does not function correctly.	ACL		08.0.01k
61779	If you configure BGP4 and route maps to redistribute static routes using the null0 interface, static routes are no longer advertised to iBGP peers at a system reload.	BGP4		08.0.01k
61821	In the web management interface, you cannot reconfigure a port interface speed to "Full Duplex".	Web Management		
68389	If you configure OSPF point-to-point links in a routing switch and then delete an OSPF neighboring router, a system reload may occur.	OSPF		08.0.01k
68891	An ACL statement that includes TCP/UDP port 0 may filter traffic incorrectly, substituting "port any" for "port 0".	ACL		08.0.01k
69334	When you configure an ACL to filter outbound IP traffic, the traffic is forwarded through the CPU.	ACL		08.0.01k

Bug ID	Bug Description	Protocol/ Feature	Version Found	Version Fixed
69765	When a corrupt header is received on an OC12 or OC48 POS module, a system reload may occur.	POS		08.0.01k
70065	If you configure multiple virtual routing interfaces on a routing switch, ARP table entries are sometimes incorrectly learned on a tagged port that does not have a matching configured subnet.	ARP		08.0.01k
70110	If OSPF is configured, entering the show ip ospf routes command may cause a system reload.	OSPF		08.0.01k
70413	If you configure a large number of ports in a VLAN, a system reload may occur.	CLI		08.0.01k

Known Issues in 08.0.01k

There are no known issues in software release 08.0.01k.

ProCurve 9300M Series Modules

Table 16 lists the modules that are currently available for use in ProCurve 9300M Series routing switches. (Discontinued modules are also listed.)

Table 16: ProCurve 9300M Series Modules

Module Type	Part Number and Description	Module String
EP Redundant Management Modules	J4885A ProCurve 9300 EP 8-Port Mini-GBIC Redundant Management Module	EP-8-port-mini-GBIC-management
EP Non-Management Modules	J4881B ProCurve 9300 EP 48-Port 10/100-TX RJ-45 Module	EP-48-port-10/100-TX-RJ45-module
	J4889B ProCurve 9300 EP 48-Port 10/100-TX Telco (RJ-21) Module	EP-48-port-10/100-TX-telco-module
	J4894A ProCurve 9300 EP 16-Port Mini-GBIC Module	EP-16-port-mini-GBIC-module
	J4895A ProCurve 9300 EP 16-Port 100/1000-T Module	EP-16-port-100/1000-T-module
	J8178A ProCurve 9300 EP 24-Port 100Base-FX Module	EP 24 Port 100Base-FX Module
Redundant Management modules (M2 and M4)	J4845A ProCurve 9300 GigLX Redundant Management Module (8-port)	8-port-gig-management-module Discontinued
	J4846A ProCurve 9300 GigSX Redundant Management Module (8-port)	8-port-gig-management-module Discontinued
	J4847A ProCurve 9300 Redundant Management Module (0-port)	0-port-management-module Discontinued

Module Type	Part Number and Description	Module String
	J4857A ProCurve 9300 Mini-GBIC Redundant Management Module (8-port)	8-port-gig-m4-management-module Discontinued
	J4879A ProCurve 9300 T-Flow Redundant Management Module (0-port)	— Discontinued
Management modules (M1) Supported only on the ProCurve 9304M and ProCurve 9308M. (M1 modules are not supported on the ProCurve 9315M.)	J4141A ProCurve 9300 10/100 Management Module (16-port)	16-port-copper-management-module Discontinued
	J4144A ProCurve 9300 Gigabit SX Management Module (8-port)	8-port-gig-management-module Discontinued
	J4146A ProCurve 9300 Gigabit 4LX/4SX Management Module (8-port)	8-port-gig-management-module Discontinued
Unmanaged Modules (Standard, non-EP)	J4140A ProCurve 9300 10/100 Module (24-port)	24-port-copper-module Discontinued
	J4142A ProCurve 9300 100Base FX Module (24-port MT-RJ)	24-port-100fx-module Discontinued
	J4143A ProCurve 9300 Gigabit SX Module (8-port)	8-port-gig-module Discontinued
	J4145A ProCurve 9300 Gigabit 4LX/4SX Module (8-port)	8-port-gig-module Discontinued
	J4842A ProCurve 9300 1000Base-T Module (8-port)	8-port-gig-copper-module Discontinued
	J4844A ProCurve 9300 GigLX Module (8-port)	8-port-gig-module Discontinued
	J4856A ProCurve 9300 Mini-GBIC Module (8-port)	8-port-gig-module Discontinued
10 Gigabit Ethernet Modules (Unmanaged, supported with both Standard and EP Management Modules)	J4891A ProCurve 9300 1-port 10 Gb Module	1-port-10Gig-module Discontinued
	J8174A ProCurve 9300 2-port 10 Gb Module	2-port-10Gig-module



© Copyright 2001, 2007 Hewlett-Packard Development Company, L.P.
The information contained herein is subject to change without notice.

5991-4755
May 2007