ProCurve Networking
HP Innovation

Release Notes:
Version E.10.23 Software
for the ProCurve Series 5300xl Switches

Release E.10.23 supports these switches:

- ProCurve Switch 5304xl (J4850A)
- ProCurve Switch 5308xl (J4819A)
- ProCurve Switch 5348xl (J4849A) – 48-port (10/100) bundle in Switch 5304xl chassis
- ProCurve Switch 5372xl (J4848A) – 72-port (10/100) bundle in Switch 5308xl chassis
- ProCurve Switch 5304xl-32G (J8166A) – 32-port (10/100/1000) bundle in 5304xl chassis
- ProCurve Switch 5308xl-48G (J8167A) – 48-port (10/100/1000) bundle in 5308xl chassis

These release notes include information on the following:

FEC, CDP Removal

Starting with Software version E.10.09, FEC trunks (Cisco Systems' FastEtherChannel for aggregated links) are no longer supported, and generation of CDP (Cisco Discovery Protocol) packets are no longer supported. In their place are IEEE standards based LACP aggregated links (as well as statically configured trunks) and generation of LLDP packets for device discovery. For more information, please see: **ftp://ftp.hp.com/pub/networking/software/LLDP-and-LACP-statement.pdf**.

Boot ROM Update Required

If your 5300xl is currently running software version E.07.37 or earlier, you must update the Boot ROM by loading and booting software version E.07.40 before installing switch software revisions later than E.07.40.

Caution

The startup-config file saved under version E.10.*xx* or greater is backward-compatible with version E.08.*xx*, but is NOT backward-compatible with E.07.*xx* or earlier software versions. Users are advised to save a copy of any pre-E.08.*xx* startup-config file BEFORE UPGRADING to E.08.*xx* or greater, in case there is ever a need to revert to pre-E.08.*xx* software. For instructions on copying the startup-config file, see Appendix A in the *Management and Configuration Guide*, available on the ProCurve Networking Web site: http:\\www.procurve.com. Click on Technical Support, then Product Manuals.

Publication Number

Part Number 5991-2127
January 2006

Applicable Product

ProCurve Switch 5304xl	(J4850A)
ProCurve Switch 5308xl	(J4819A)
ProCurve Switch 5348xl	(J4849A)
ProCurve Switch 5372xl	(J4848A)
ProCurve Switch 5304xl-32G	(J8166A)
ProCurve Switch 5308xl-48G	(J8167A)

Trademark Credits

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

Software Credits

SSH on ProCurve Switches is based on the OpenSSH software toolkit. This product includes software developed by the OpenSSH Project for use in the OpenSSH Toolkit. For more information on OpenSSH, visit

http:// www.openssh.com.

SSL on ProCurve Switches is based on the OpenSSL software toolkit. This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit. For more information on OpenSSL, visit

http://www.openssl.org.

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

Disclaimer

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

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 Boulevard, m/s 5551
Roseville, California 95747-5551
www.procurve.com

Contents

Software Management

Clarifications

Enhancements

Software Fixes in Release E.06.xx through E.10.xx

Software Management

Software Updates

Check the ProCurve Networking Web site frequently for free software updates for the various ProCurve switches you may have in your network.

Downloading Switch Documentation and Software from the Web

You can download software updates and the corresponding product documentation from HP's ProCurve web site as described below.

To Download a Software Version:

1. Go to the ProCurve Networking Web site at:

   http://www.procurve.com.

2. Click on **Software updates** (in the sidebar).
3. Under **Latest software**, click on **Switches**.

To Download Product Documentation: You will need the Adobe® Acrobat® Reader to view, print, and/or copy the product documentation.

1. Go to HP's ProCurve web site at **http://www.procurve.com**.
2. Click on **Technical support**, then **Product manuals**.
3. Click on the name of the product for which you want documentation.
4. On the resulting web page, double-click on a document you want.
5. When the document file opens, click on the disk icon  in the Acrobat® toolbar and save a copy of the file.

Downloading Software to the Switch

Caution

The startup-config file generated by the latest software release may not be backward-compatible with the same file generated by earlier software releases. Refer to "Boot ROM Update Required" on the front page.

HP periodically provides switch software updates through the ProCurve Networking Web site (**http://www.procurve.com**). After you acquire the new software file, you can use one of the following methods for downloading it to the switch:

- For a TFTP transfer from a server, do either of the following:
  - Click on **Download OS** in the Main Menu of the switch's menu interface and use the (default) **TFTP** option.
  - Use the **copy tftp** command in the switch's CLI (see below).
- For an Xmodem transfer from a PC or Unix workstation, do either of the following:
  - Click on **Download OS** in the Main Menu of the switch's menu interface and select the **Xmodem** option.
  - Use the `copy xmodem` command in the switch's CLI (page 4).
- Use the download utility in ProCurve Manager Plus.
- A switch-to-switch file transfer

Note

Downloading new software does not change the current switch configuration. The switch configuration is contained in a separate file that can also be transferred, for example, for archive purposes or to be used in another switch of the same model.

This section describes how to use the CLI to download software to the switch. You can also use the menu interface for software downloads. For more information, refer to the *Management and Configuration Guide* for your switch.

TFTP Download from a Server

Syntax: **copy tftp flash <ip-address> <remote-os-file> [< primary | secondary >]**

Note that if you do not specify the flash destination, the TFTP download defaults to the primary flash.

For example, to download a software file named E_10_2x.swi from a TFTP server with the IP address of 10.28.227.103:

1. Execute the copy command as shown below:

   ```
   HPswitch # copy tftp flash 10.28.227.103 E_10_2x.swi
   The primary OS image will be deleted. continue [y/n]? Y
   03125K
   ```

2. When the switch finishes downloading the software file from the server, it displays the progress message:

   ```
   Validating and Writing System Software to FLASH..
   ```

   When the CLI prompt re-appears, the switch is ready to reboot to activate the downloaded software

3. Reboot the switch.

   After the switch reboots, it displays the CLI or Main Menu, depending on the **Logon Default** setting last configured in the menu's Switch Setup screen.

Xmodem Download From a PC or Unix Workstation

This procedure assumes that:

- The switch is connected via the Console RS-232 port on a PC operating as a terminal. (Refer to the Installation Guide you received with the switch for information on connecting a PC as a terminal and running the switch console interface.)
- The switch software is stored on a disk drive in the PC.
- The terminal emulator you are using includes the Xmodem binary transfer feature. (For example, in the Microsoft Windows NT® terminal emulator, you would use the **Send File** option in the **Transfer** drop-down menu.)

Syntax: **copy xmodem flash < unix | pc >**

For example, to download a software file from a PC:

1. To reduce the download time, you may want to increase the baud rate in your terminal emulator and in the switch to a value such as 57600 bits per second. (The baud rate must be the same in both devices.) For example, to change the baud rate in the switch to 57600, execute this command:

   ```
   HPswitch(config)# console baud-rate 57600
   ```

   (If you use this option, be sure to set your terminal emulator to the same baud rate.)

2. Execute the following command in the CLI:

   ```
   ProCurve # copy xmodem flash primary
   The primary OS image will be deleted. continue [y/n]? Y
   Press 'enter' and start XMODEM on your host . . .
   ```

3. Execute the terminal emulator commands to begin the Xmodem transfer. The download can take several minutes, depending on the baud rate used in the transfer.
4. Reboot the switch.

   After the switch reboots, it displays the CLI or Main Menu, depending on the **Logon Default** setting last configured in the menu's Switch Setup screen.

5. Use the following command to confirm that the software downloaded correctly:

   HPswitch> show system

   Check the **Firmware revision** line to verify that the switch cownloaded the new software.

6. If you increased the baud rate on the switch (step 1), use the same command to return it to its previous setting. (HP recommends a baud rate of 9600 bits per second for most applications.) Remeber to return your terminal emulator to the same baud rate as the switch.

Saving Configurations While Using the CLI

The switch operates with two configuration files:

- **Running-Config File:** Exists in volatile memory and controls switch operation. Rebooting the switch erases the current running-config file and replaces it with an exact copy of the current startup-config file. To save a configuration change, you must save the running configuration to the startup-config file.
- **Startup-Config File:** Exists in flash (non-volatile) memory and preserves the most recently-saved configuration as the "permanent" configuration. When the switch reboots for any reason, an exact copy of the current startup-config file becomes the new running-config file in volatile memory.

When you use the CLI to make a configuration change, the switch places the change in the running-config file. If you want to preserve the change across reboots, you must save the change to the startup-config file. Otherwise, the next time the switch reboots, the change will be lost. There are two ways to save configuration changes while using the CLI:

- Execute **write memory** from the Manager, Global, or Context configuration level.
- When exiting from the CLI to the Main Menu, press **[Y]** (for `Yes`) when you see the "save configuration" prompt:

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

ProCurve Switch Software Key

Software Letter	ProCurve Switch, Routing Switch, or Router
C	1600M, 2400M, 2424M, 4000M, and 8000M
CY	Switch 8100fl Series (8108fl and 8116fl)
E	Switch 5300xl Series (5304xl, 5308xl, 5348xl, and 5372xl)
F	Switch 2500 Series (2512 and 2524), Switch 2312, and Switch 2324
G	Switch 4100gl Series (4104gl, 4108gl, and 4148gl)
H	Switch 2600 Series, Switch 2600-PWR Series: H.07.81 and earlier, or H.08.55 and greater, Switch 2600-8-PWR requires H.08.80 or greater.
H	Switch 6108: H.07.xx and earlier
I	Switch 2800 Series (2824 and 2848)
J	Secure Router 7000dl Series (7102dl and 7203dl)
M	Switch 3400cl Series (3400-24G and 3400-48G) and Series 6400cl (CX4 6400cl-6XG and X2 6400cl-6XG)
N/A	Switch 9408sl, 9300 Series (9304M, 9308M, and 9315M), Switch 6208M-SX and Switch 6308M-SX (Uses software version number only; no alphabetic prefix. For example 07.6.04.)

Minimum Software Versions for Series 5300xl Switch Features

For Software Features. To view a tabular listing of major switch software features and the minimum software version each feature requires:

1. Visit the ProCurve Networking Web site at **http://www.procurve.com**.
2. Click on **software.**
3. Click on **Minimum Software Version Required by Feature**.

If you are viewing this publication online, just click on step 3 to go directly to the Minimum Software information.

For Switch 5300xl Hardware Accessories.

ProCurve Device	Minimum Supported Software Version
J4820A 24-Port 10/100-TX Module	E.05.04
J4821A 4-Port 100/1000-T Module	E.05.04
J4839A Redundant Power Supply (RPS)	E.05.04
J4852A 12-Port 100-FX MTRJ Module	E.06.10
J4878A 4-Port Mini-GBIC Module	E.05.04
J4858A Gigabit-SX-LC Mini-GBIC	E.05.04
J4859A Gigabit-LX-LC Mini-GBIC	E.05.04
J4860A Gigabit-LH-LC Mini-GBIC	E.06.01
J8161A 24-Port 10/100-TX PoE Module	E.08.22
J4907A 16-Port 10/100/1000-T Module	E.08.42
J8162A Access Controller xl Module	E.09.21
J8177B 1000Base-T Mini-GBIC	E.09.22

Clarifications

LLDP and LACP

Starting with Software version E.10.10, FEC trunks (Cisco Systems' FastEtherChannel for aggregated links) are no longer supported, and generation of CDP (Cisco Discovery Protocol) packets are no longer supported. In their place are IEEE standards-based LACP aggregated links (as well as statically configured trunks) and generation of LLDP packets for device discovery.

For more information, please see: **ftp://ftp.hp.com/pub/networking/software/LLDP-and-LACP-statement.pdf**.

Mesh Design Optimization

Mesh performance can be enhanced by using mesh designs that are as small and compact as possible while still meeting the network design requirements. The following are limits on the design of meshes and have not changed:

1. Any switch in the mesh can have up to 24 meshed ports.
2. A mesh domain can contain up to 12 switches.
3. Up to 5 inter-switch meshed hops are allowed in the path connecting two nodes.
4. A fully interconnected mesh domain can contain up to 5 switches.

Mesh performance can be optimized by keeping the number of switches and the number of possible paths between any two nodes as small as possible. As mesh complexity grows, the overhead associated with dynamically calculating and updating the cost of all of the possible paths between nodes grows exponentially. Cost discovery packets are sent out by each switch in the mesh every 30 seconds and are flooded to all mesh ports. Return packets include a cost metric based on inbound and outbound queue depth, port speed, number of dropped packets, etc. Also, as mesh complexity grows, the number of hops over which a downed link has to be reported may increase, thereby increasing the reconvergence time.

The simplest design is the two-tier design because the number of possible paths between any two nodes is kept low and any bad link would have to be communicated only to it's neighbor switch.

As shown here, meshing allows multiple redundant links between switches in the domain, and more than one link is allowed between any two switches in the domain.

Note also that a switch can have up to 24 ports configured for meshing.

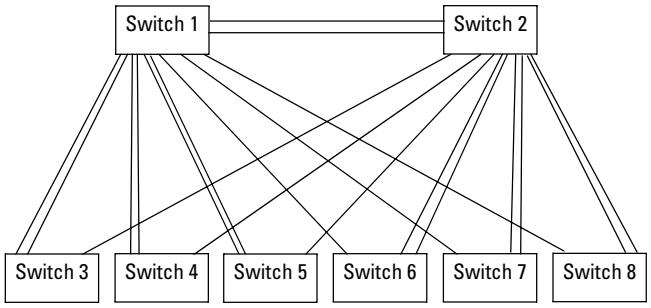

Figure 1. Example of a Two-Tier Mesh Design

Other factors affecting the performance of mesh networks include the number of destination addresses that have to be maintained, and the overall traffic levels and patterns. However a conservative approach when designing new mesh implementations is to use the two-tier design and limit the mesh domain to eight switches where possible.

For more information, refer to the chapter titled “Switch Meshing” in the Advanced Traffic Management Guide for your switch.

General Switch Traffic Security Guideline

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

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

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

The Management VLAN IP Address

The optional Management VLAN, if used, must be configured with a manual IP address. It does not operate with DHCP/Bootp configured for the IP address.

Heavy Memory Usage with PIM-DM

Heavy use of PIM (Many S/G--source-group--flows over many VLANs) combined with other memory-intensive features, can oversubscribe memory resources and impact overall performance. If available memory is exceeded, the switch drops any new multicast flows, and generates appropriate log messages. Corrective actions can include reducing the number of VLANs on the 5300xl device by moving some VLANs to another device, free up system resources by disabling another, non-PIM feature, and/or moving some hosts to another device. For more information, refer to "Operating Notes" and "Messages Related to PIM Operation" in the chapter titled "PIM DM (Dense Mode)" in the *Advanced Traffic Management Guide* (February, 2004 or later) for the ProCurve Series 5300xl switches. For more information on PIM-DM operation, refer to the chapter titled "PIM-DM (Dense Mode)" in the *Advanced Traffic Management Guide* for the ProCurve Series 5300xl switches. (To download switch documentation for software release E.09.*xx*, refer to "Software Updates" on page 1.)

Change in QoS Priority and Policy Limit

Beginning with software release E.09.22, the switch allows configuration of up to 250 priority and/or DSCP policy configurations. Attempting to add more than 250 entries generates an error message in the CLI.

Heavy use of QoS, combined with other memory-intensive features, can oversubscribe memory resources and impact overall performance. Updating the switch software from an earlier release in which more than 250 entries were configured causes the switch to drop any entries in excess of the first 250 and to generate an event log message indicating this action. For more information, refer to "QoS Operating Notes" in the chapter titled "Quality of Service (QoS): Managing Bandwidth More Effectively" in the *Advanced Traffic Management Guide* for the ProCurve Series 5300xl switches (part number 5990-6051, January 2005 or later). Note that the above limit supercedes the limit indicated in the January 2005 edition of the *Advanced Traffic Management Guide*. To download switch documentation for software release E.09.22, refer to "Software Updates" on page 1.

Interoperating with 802.1s Multiple Spanning-Tree

The ProCurve implementation of Multiple Spanning-Tree (MSTP) in software release E.08.*xx* and greater complies with the IEEE 802.1s standard and interoperates with other devices running compliant versions of 802.1s. Note that the ProCurve Series 9300 routing switches do not offer 802.1s-compliant MSTP. Thus, to support a connection between a 9300 routing switch and a 5300xl switch running MSTP, configure the 9300 with either 802.1D (STP) or 802.1w (RSTP). For more information

on this topic, refer to the chapter titled "Spanning-Tree Operation" in the *Advanced Traffic Management Guide* (part number 5990-6051, January 2005 or later). (To download switch documentation for software release E.09.22, refer to "Software Updates" on page 1.)

Rate-Limiting

The configured rate limit on a port reflects the permitted forwarding rate from the port to the switch backplane, and is visible as the *average* rate of the outbound traffic originating from the rate-limited port. (The most accurate rate-limiting is achieved when using standard 64-byte packet sizes.) Also, rate-limiting reflects the available percentage of a port's entire inbound bandwidth. The rate of inbound flow for traffic of a given priority and the rate of flow from a rate-limited port to a particular queue of an outbound port are not measures of the actual rate limit enforced on a port. Also, rate-limiting is byte-based and is applied to the available bandwidth on a port, and not to any specific applications running through the port. If the total bandwidth requested by all applications together is less than the available, configured maximum rate, then no rate-limit can be applied. This situation occurs with a number of popular throughput-testing software applications, as well as most regular network applications.

As a performance consideration, implementing rate-limiting in heavy traffic situations involving QoS, can affect overall performance. For more information on rate-limiting operation, refer to "Operating Notes for Rate-Limiting" in the chapter titled "Optimizing Traffic Flow with Port Controls, Port Trunking, and Filters" of the *Management and Configuration Guide* (part number 5990-6050, January 2005 or later) for the ProCurve Series 5300xl switches. (To download switch documentation for software release E.09.22, refer to "Software Updates" on page 1.)

OS/Web/Java Compatibility Table

The switch web agent supports the following combinations of OS browsers and Java Virtual Machines:

Operating System	Internet Explorer	Java
Windows NT 4.0 SP6a	5.00, 5.01 5.01, SP1 6.0, SP1	Sun Java 2 Runtime Environment: – Version 1.3.1.12 – Version 1.4.2.05
Windows 2000 Pro SP4	5.05, SP2 6.0, SP1	
Windows XP Pro SP2	6.0, SP1	Sun Java 2 Runtime Environment: – Version 1.5.0.02
Windows Server SE 2003 SP1	6.0, SP1	

Time Zone Offset

Starting with release E.05.*xx*, the method of configuring the Time Zone for TimeP or SNTP configuration has been updated. Previous switch software for all ProCurve switches used positive time offset values for time zones that are West of GMT and negative values for time zones that are East of GMT. The standards indicate that time zones West of GMT should be designated by negative offset values, and time zones East of GMT by positive values. Software version E.05.*xx* updates this configuration method, but if you use the same values for indicating time zones as you did for previous ProCurve switches, the time will be set incorrectly on your Series 5300GL switch. For example, for previous ProCurve switches, the US Pacific time zone was configured by entering **+480**. With software version E.05.*xx*, the US Pacific time zone must now be configured by entering **-480**.

Enhancements

Unless otherwise noted, each new release includes the features added in all previous releases.

Release E.10.03 through E.10.20 Enhancements

Software fixes only; no new enhancements. Versions E.10.11 to E.10.19 were never built.

Release E.10.02 Enhancements (June 2005)

The table below summarizes enhancements made in this release. New features are described in detail in the sections following the table.

Enhancement	Overview
Custom Login Banners for Console and Web Browser Access	Configurable login banners are now available for console (CLI and Menu interfaces) and the Web browser interface (page 13).
ECMP (Equal-Cost Multi-path) for Different Subnets Through the Same Next-Hop Routes	5300xl routers now support optional load-sharing across redundant links where the network offers two, three, or four equal-cost next-hop routes for traffic to different subnets. (page 18).
Disable TFTP	This feature enables increased security by automatically disabling TFTP and auto-TFTP when secure FTP (SFTP) is invoked, or allowing a system operator using secure FTP (SFTP) to disable TFTP operation. This feature also enables a system operator to manually disable TFTP and/or auto-TFTP (page 20).
Change in the Rate-Limiting Command	The syntax of the rate-limiting command included in the switch software has changed to accomodate the new ICMP rate-limiting feature available with E.10.02 (page 22).
ICMP Rate-Limiting	ICMP Rate-Limiting allows a system operator to restrict ICMP traffic to levels that permit necessary ICMP functions, but throttle additional traffic that may be due to worms or viruses (reducing their spread and effect). In addition, this preserves inbound port bandwidth for non-ICMP traffic (page 23).
LLDP-MED	LLDP Media Endpoint Discovery provides an extension to LLDP (Link Layer Discover Protocol) to support VoIP deployments. Information on this feature is provided in a separate publication titled *Manual Supplement: LLDP and LLDP-MED for the HP ProCurve Series 5300xl Switches* on the ProCurve website at www.procurve.com.
RADIUS Accounting Enhancement	The RADIUS accounting features available for 802.1X access to the switch are now available when using Web Authentication and MAC authentication for switch access. (page 32).
RADIUS-Assigned ACLs	This feature uses RADIUS-assigned, per-port ACLs for Layer-3 filtering of inbound IP traffic from authenticated clients (page 33).

Custom Login Banners for the Console and Web Browser Interfaces

You can now configure the switch to display a login banner of up to 320 characters when an operator initiates a management session with the switch through any of the following methods:

- Telnet
- serial connection
- SSHv2 (SSHv1 does not include support for banners.)
- Web browser

In the factory default configuration, the switch displays the following default banner:

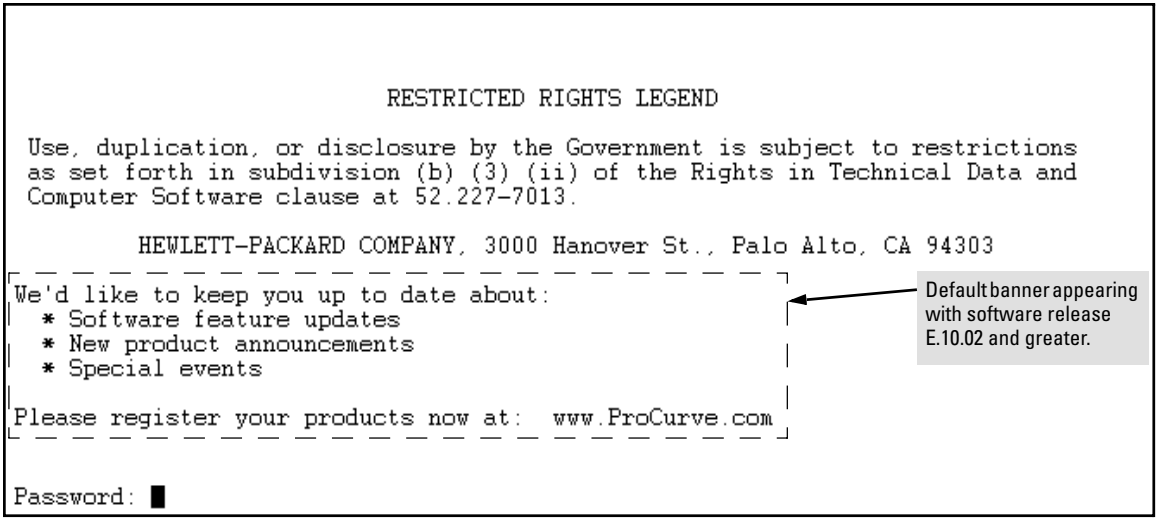

Figure 2. The Default Login Banner

Note

The switch's Web browser interface does not display the default banner.

Banner Operation with Telnet, Serial, or SSHv2 Access

When a system operator begins a login session, the switch displays the banner above the local password prompt or, if no password is configured, above the **Press any key to continue prompt.** Entering a correct password or, if no password is configured, pressing any key clears the banner from the CLI and displays the CLI prompt. (Refer to figure 2 on page 13.)

Banner Operation with Web Browser Access

When a system operator uses a Web broswer to access the switch, the text of a non-default banner configured on the switch appears in a dedicated banner window with a link to the Web agent home page. Clicking on **To Home Page** clears the banner window and prompts the user for a password (if configured). Following entry of the correct username/password information (or if no username/password is required), the switch then displays either the Registration page or the switch's home page. Note that if the banner feature is disabled or if the switch is using the factory-default banner shown in figure 2, then the banner page does not appear in the Web browser when an operator initiates a login session with the switch.

Configuring and Displaying a Non-Default Banner

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

1. Enable non-default banner operation and define the endpoint delimiter for the banner.
2. Enter the desired banner text, including any specific line breaks you want.
3. Enter the endpoint delimiter.
4. Use **show banner motd** to display the current banner status.

Syntax: banner motd < *delimiter* >
no banner motd

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

< *banner-text-string* >

*The switch allows up to 320 banner characters, including blank spaces and CR-LF (***[Enter]***). (The tilde " **~**" and the delimiter defined by* **banner motd <*delimiter*>** *are not allowed as part of the banner text.) While entering banner text, you can backspace to edit the current line (that is, a line that has not been terminated by a CR-LF.) However, terminating a line in a banner by entering a CR-LF prevents any further editing of that line. To edit a line in a banner entry after terminating the line with a CR-LF requires entering the delimiter described above and then re-configuring new banner text.*

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

Example of Configuring and Displaying a Banner

Suppose a system operator wanted to configure the following banner message on her company's 5300xl switches:

```
This is a private system maintained by the

         Allied Widget Corporation.

Unauthorized use of this system can result in

        civil and criminal penalties!
```

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

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

Figure 3. Example of Configuring a Login Banner

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

```
ProCurve(config)# show banner motd

Banner Information

Banner status: Enabled
Configured Banner:

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

Figure 4. Example of show banner motd Output

```
ProCurve(config)# show running

Running configuration:

; J4850A Configuration Editor; Created on release #E.10.02

hostname "ProCurve"
module 1 type J8161A
module 2 type J8161A
snmp-server community "notpublic" Unrestricted
vlan 1
   name "DEFAULT_VLAN"
   untagged A1-A24,B1-B24
   ip address dhcp-bootp
   exit
banner motd "    This is a private system maintained by the
             Allied Widget Corporation.
    Unauthorized use of this system can result in
          civil and criminal penalties!"
password manager
password operator
```

Shows the current banner configuration.

Figure 5. The Current Banner Appears in the Switch's Running-Config File

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

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

                        RESTRICTED RIGHTS LEGEND

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

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

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

Password:
```

The login screen displays the configured banner.

Entering a correct password clears the banner and displays the CLI prompt.

Figure 6. Example of CLI Result of the Login Banner Configuration

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

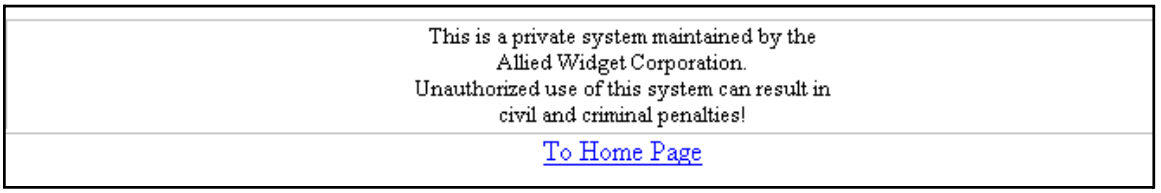

Figure 7. Example of Web Browser Interface Result of the Login Banner Configuration

Operating Notes

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

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

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

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

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

  Notice to all users

OSPF Equal-Cost Multipath (ECMP) for Different Subnets Available Through the Same Next-Hop Routes

Using software prior to release E.10.*xx*, if different subnet destinations in an OSPF network are reachable through a set of equal-cost next-hop routes, the router chooses the same next-hop route for traffic to all of these destinations. Beginning with software release E.10.*xx*, 5300xl routers support optional load-sharing across redundant links where the network offers two, three, or four equal-cost next-hop routes for traffic to different subnets. (All traffic for different hosts in the same subnet goes through the same next-hop router.)

For example, in the OSPF network shown below, IP load-sharing is enabled on router "A". In this case, OSPF calculates three equal-cost next-hop routes for each of the subnets and then distributes per-subnet route assignments across these three routes.

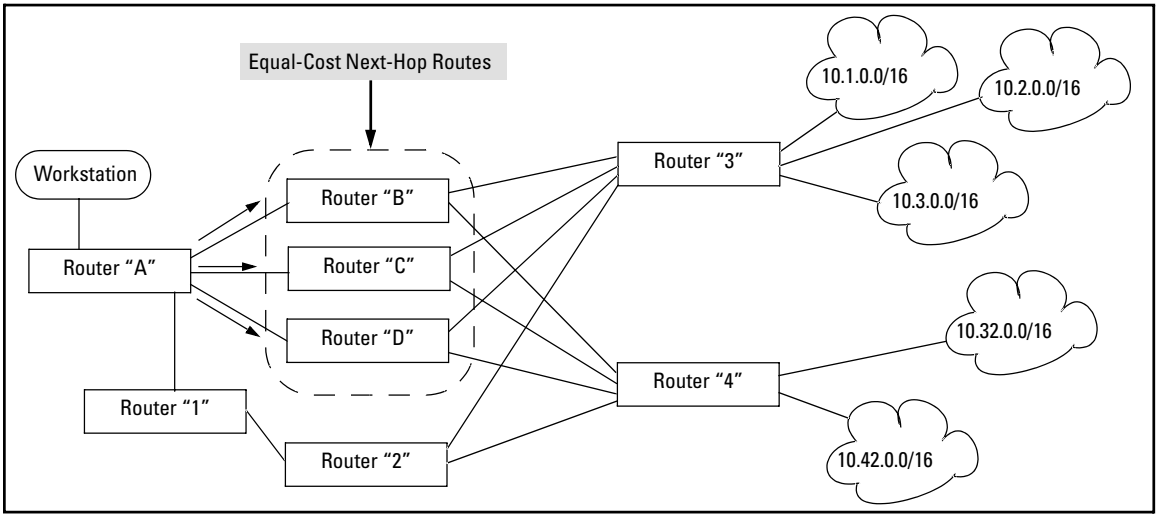

Figure 8. Example of Load-Sharing Traffic to Different Subnets Through Equal-Cost Next-Hop Routers

Table 1. Example of a Routing Table for the Network in Figure 8

Destination Subnet	Router "A" Next Hop
10.1.0.0/16	Router "C"
10.2.0.0/16	Router "D"
10.3.0.0/16	Router "B"
10.32.0.0/16	Router "B"
10.42.0.0/16	Router "D"

Note that IP load-sharing does not affect routed traffic to different hosts on the same subnet. That is, all traffic for different hosts on the same subnet will go through the same next-hop router. For

example, if subnet 10.32.0.0 includes two servers at 10.32.0.11 and 10.32.0.22, then all traffic from router "A" to these servers will go through router "B".

Syntax: [no] ip load-sharing < 2 - 4 >

When OSPF is enabled and multiple, equal-cost, next-hop routes are available for traffic destinations on different subnets, this feature, by default, enables load-sharing among up to four next-hop routes. The **no** *form of the command disables this load-sharing so that only one route in a group of multiple, equal-cost, next-hop routes is used for traffic that could otherwise be load-shared across multiple routes. For example, in figure 8 on page 18, the next-hop routers "B", "C", and "D" are available for equal-cost load-sharing of eligible traffic. Disabling IP load-sharing means that router "A" selects only one next-hop router for traffic that is actually eligible for load-sharing through different next-hop routers. (Default: Enabled with four equal-cost, next-hop routes allowed)*

Note: *In the default configuration, IP load-sharing is enabled by default. However, it has no effect unless IP routing and OSPF are enabled.*

< 1 - 4 >

Specifies the maximum number of equal-cost next hop paths the router allows. (Range: **2** *-* **4***; Default:* **4***)*

Displaying the Current IP Load-Sharing Configuration

Use the **show running** command to view the currently active IP load-sharing configuration, and **show config** to view the IP load-sharing configuration in the startup-config file. (While in its default configuration, IP load-sharing does not appear in the command output.) If IP load sharing is configured with non-default settings (disabled or configured for either two or three equal-cost next-hop paths), then the current settings are displayed in the command output.

```
HP ProCurve Switch 5304XL(config)# show running

Running configuration:

; J4850A Configuration Editor; Created on
release #E.10.00

hostname "HP ProCurve Switch 5304XL"
module 1 type J4820A
snmp-server community "public" Unrestricted
vlan 1
   name "DEFAULT_VLAN"
   untagged A1-A24
   ip address dhcp-bootp
   exit
ip load-sharing 3
access-controller vlan-base 2000
```

Indicates a non-default IP load-sharing configuration allowing three equal-cost next-hop paths for routed traffic with different subnet destinations. If the router is configured with the default IP load-sharing configuration, IP load-sharing does not appear in the **show config** or **show running** command output.

Figure 9. Displaying a Non-Default IP Load-Sharing Configuration

Disable TFTP and Auto-TFTP for Enhanced Security

Beginning with software release E.10.02, using the **ip ssh filetransfer** command to enable Secure FTP (SFTP) automatically disables TFTP and auto-TFTP (if either or both are enabled).

```
ProCurve(config)# ip ssh filetransfer
Tftp and auto-tftp have been disabled.
ProCurve(config)# sho run

Running configuration:

; J4850A Configuration Editor; Created on release #E.10.02

hostname "ProCurve"
module 1 type J8161A
module 2 type J8161A
vlan 1
   name "DEFAULT_VLAN"
   untagged A1-A24,B1-B24
   ip address 10.28.234.176 255.255.240.0
   exit
ip ssh filetransfer
no tftp-enable
password manager
password operator
```

Enabling SFTP automatically disables TFTP and auto-tftp and displays this message.

Viewing the configuration shows that SFTP is enabled and TFTP is disabled.

Figure 10. Example of Switch Configuration with SFTP Enabled

If you enable SFTP, then later disable it, TFTP and auto-TFTP remain disabled unless they are explicitly re-enabled.

Operating Rules

- The TFTP feature is enabled by default, and can be enabled or disabled through the CLI, the Menu interface, or an SNMP application. Auto-TFTP is disabled by default and must be configured through the CLI.

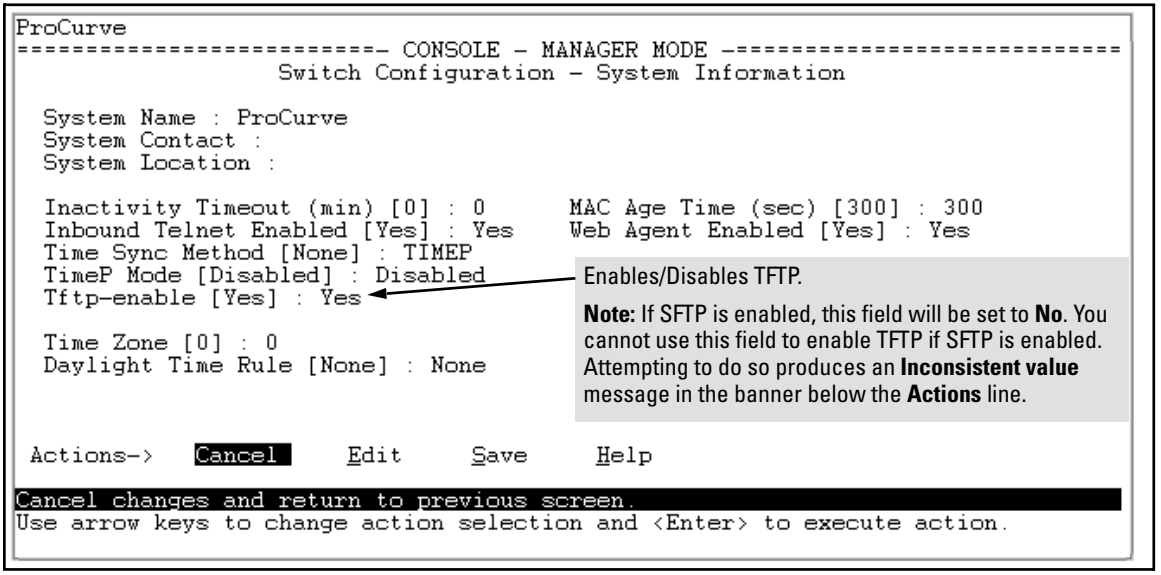

Figure 11. Using the Menu Interface To Disable TFTP

- While SFTP is enabled, TFTP and auto-TFTP cannot be enabled from the CLI. Attempting to enable either non-secure TFTP option while SFTP is enabled produces one of the following messages in the CLI:

  ```
  SFTP must be disabled before enabling tftp.

  SFTP must be disabled before enabling auto-tftp.
  ```

  Similarly, while SFTP is enabled, TFTP cannot be enabled using an SNMP management application. Attempting to do so generates an "inconsistent value" message. (An SNMP management application cannot be used to enable or disable auto-TFTP.)

- To enable SFTP by using an SNMP management application, you must first disable TFTP and, if configured, auto-TFTP on the switch. You can use either an SNMP application or the CLI to disable TFTP, but must use the CLI to disable auto-TFTP. The following two CLI commands disable TFTP and auto-TFTP on the switch.

Syntax: no tftp-enable

This command disables all TFTP operation on the switch except for the auto-TFTP feature. To re-enable TFTP operation, use the **tftp-enable** *command. When TFTP is disabled, the instances of* **tftp** *in the CLI copy command and the Menu interface "Download OS" screen become unavailable.*

Note: *This command does* ***not*** *disable auto-TFTP operation. To disable an auto-TFTP command configured on the switch, use the* **no auto-tftp** *command described below to remove the command entry from the switch's configuration.*

Syntax: no auto-tftp

If auto-TFTP is configured on the switch, this command deletes the **auto-tftp** *entry from the switch configuration, thus preventing auto-tftp operation if the switch reboots.*

Note: *This command does not affect the current TFTP-enable configuration on the switch.*

Change in the Rate-Limiting CLI Command

Beginning with software release E.10.02, the syntax of the rate-limiting command included in the switch software has changed to accomodate the new ICMP rate-limiting feature available with E.10.02 (page 23).

Syntax: [no] int < *port-list* > rate-limit < 0..100 >

This command syntax applies to software releases prior to E.10.xx, and applies to inbound rate-limiting of all traffic received on a specific port.

Syntax: [no] int < *port-list* > rate-limit < all | icmp > < 0..100 >

This command syntax replaces the above rate-limit command and includes options for rate-limiting of either all inbound traffic on a port or just the inbound ICMP traffic on a port.

- *For more information on rate-limiting all traffic inbound on a port, refer to the section titled "Rate-Limiting" in the chapter titled "Port Traffic Controls" in the Management and Configuration Guide for your switch.*
- *For information on rate-limiting only the inbound ICMP traffic on a port, refer to "ICMP Rate-Limiting" on page 23 in this publication.*

ICMP Rate-Limiting

In IP networks, ICMP messages are generated in response to either inquiries or requests from routing and diagnostic functions. These messages are directed to the applications originating the inquiries. In unusual situations, if the messages are generated rapidly with the intent of overloading network circuits, they can threaten network availability. This problem is visible in denial-of-service (DoS) attacks or other malicious behaviors where a worm or virus overloads the network with ICMP messages to an extent where no other traffic can get through. (ICMP messages themselves can also be misused as virus carriers). Such malicious misuses of ICMP can include a high number of ping packets that mimic a valid source IP address and an invalid destination IP address (spoofed pings), and a high number of response messages (such as Destination Unreachable error messages) generated by the network. ICMP Rate-Limiting provides a method for limiting the amount of bandwidth that may be utilized for inbound ICMP traffic on a switch port or trunk. This feature allows users to restrict ICMP traffic to levels that permit necessary ICMP functions, but throttle additional traffic that may be due to worms or viruses (reducing their spread and effect). In addition, this preserves inbound port bandwidth for non-ICMP traffic.

Terminology

All-Traffic Rate-Limiting: Applies a rate-limit to all inbound traffic, including ICMP traffic, received on an interface.

ICMP Rate-Limiting: Applies a rate-limit to all inbound ICMP traffic received on an interface, but does not limit other types of inbound traffic.

Spoofed Ping: An ICMP echo request packet intentionally generated with a valid source IP address and an invalid destination IP address. Spoofed pings are often created with the intent to oversubscribe network resources with traffic having invalid destinations.

Effect of ICMP Rate-Limiting

ICMP rate-limiting generally allows only a specified percentage of an interface's inbound bandwidth to be used for ICMP traffic. As a result, inbound bandwidth is preserved for non-ICMP traffic and the port or trunk throttles any sudden flood of inbound ICMP traffic that may be due to a worm or virus attack (or any other cause). Notice that ICMP rate-limiting does not throttle non-ICMP traffic. In cases where you want to throttle both ICMP traffic and all other inbound traffic on a given interface, you can configure both ICMP rate-limiting and all-traffic rate-limiting.

Caution

The ICMP protocol is necessary for routing, diagnostic, and error responses in an IP network. ICMP rate-limiting is primarily used for throttling worm or virus-like behavior, and should normally be configured to allow one to five per cent of available inbound bandwidth to be used for ICMP traffic. ***This feature should not be used to remove all ICMP traffic from a network.***

Note

Because all-traffic rate-limiting and ICMP rate-limiting operate similarly, the CLI command for the all-traffic version of rate-limiting has been modified for compatibility with the ICMP rate-limiting CLI command. Beginning with software release E.10.02, these commands appear in the following format:

rate-limit [all | icmp]

For more on all-traffic rate-limiting, refer to the chapter titled "Port Traffic Controls" in the *Management and Configuration Guide* for your switch (January 2005 or greater).

Network Application. Apply ICMP rate-limiting on all connected interfaces on the switch to effectively throttle excessive ICMP messaging from any source. On edge interfaces, where ICMP traffic should be minimal, a threshold of 1% of available bandwidth should be sufficient for most applications. On core interfaces, such as switch-to-switch and switch-to-router, a maximum threshold of 5% should be sufficient for normal ICMP traffic. ("Normal" ICMP traffic levels should be the maximums that occur when the network is rebooting.)

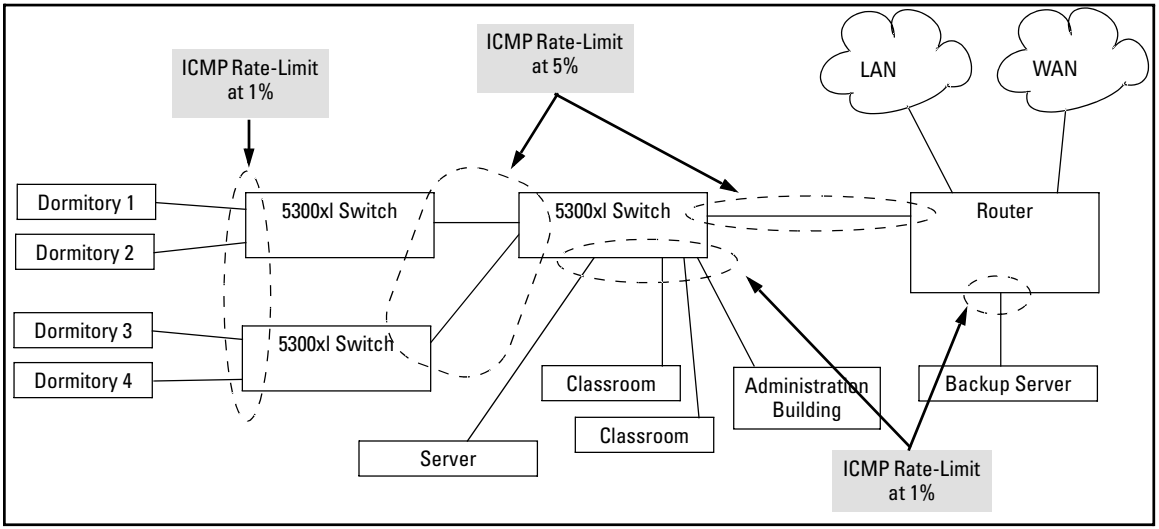

Figure 12. Example of ICMP Rate-Limiting

ICMP Rate-Limiting Operation. ICMP rate-limiting operates on an interface (per-port or per-trunk) basis to allow, on average, the highest expected amount of legitimate, inbound ICMP traffic. For example, if a 100 Mbps port negotiates a link to another switch at 100 Mbps and is ICMP rate-limit configured at 5%, then the inbound ICMP traffic flow through that port is limited to 5 Mbps. Similarly, if the same port negotiates a 10 Mbps link, then it allows 0.5 Mbps of inbound traffic. (For more on performance under varying operating conditions, refer to "Operating Notes for ICMP Rate-Limiting" on page 29.) If an interface experiences an inbound flow of ICMP traffic in excess of its configured limit, the switch generates a log message and an SNMP trap (if an SNMP trap receiver is configured).

Applying ICMP Rate-Limiting to a Port Trunk. These rules apply to ICMP rate-limiting when applied to a trunk:

- The configured ICMP traffic limit is applied as a percentage of all traffic inbound on the trunk.
- ICMP rate-limiting is only supported on a port trunk where all members of the trunk are *in the same module slot*. ICMP rate-limiting is **not** supported on trunks having members in multiple module slots.
- All ports belonging to a trunk configured for ICMP rate-limiting operate according to the trunk configuration, regardless of the ICMP rate-limiting state that existed on the port prior to its being added to the trunk. (While a port is in a trunk, any ICMP rate-limiting previously configured for that port is suspended, but remains in the switch configuration.)
- Removing a port from a trunk returns the port to whatever ICMP rate-limiting state existed on the port before it was put into the trunk.

Note

A rate-limited trunk should include only ports on the same slot/module. A rate-limited trunk configured across module boundaries is not supported and produces unpredictable rate-limiting operation and results.

Using Both ICMP and All-Traffic Rate-Limiting on an Interface. ICMP and all-traffic rate-limiting can be configured on the same interface. All-Traffic rate-limiting applies to all inbound traffic(including ICMP traffic), while ICMP rate-limiting applies only to inbound ICMP traffic.

Note that if the inbound, all-traffic load on an interface meets or exceeds the current all-traffic rate-limit while the ICMP traffic rate-limit on the same interface has not been reached, then all excess traffic will be dropped, including any inbound ICMP traffic above the all-traffic limit (regardless of whether the ICMP rate-limit has has been reached). Suppose, for example:

- The all-traffic limit on port “X” is configured at 55% of the port’s bandwidth.
- The ICMP traffic limit on port “X” is configured at 2% of the port’s bandwidth.

If at a given moment:

- inbound ICMP traffic on port “X” is using 1% of the port’s bandwidth, and
- inbound traffic of all types on port “X” demands 61% of the ports’s bandwidth,

then all inbound traffic above 55% of the port’s bandwidth, including any additional ICMP traffic will be dropped as long as all inbound traffic combined on the port demands 55% or more of the port’s bandwidth.

Note

Under network stress conditions, an interface may allow occasional bursts of inbound ICMP traffic forwarding that exceed the interface’s configured ICMP traffic rate. Refer to “ICMP Rate-Limit Imposes an Average Bandwidth Limit” on page 30.

Configuring Inbound Rate-Limiting. This command controls inbound usage of a port by setting a limit on the bandwidth available for inbound traffic.

Syntax: [no] int < *port- list* | *trunk-list* > rate-limit icmp < 0..100 >

Configures inbound ICMP traffic rate limiting. You can configure a rate limit from either the global configuration level (as shown above) or from the interface context level. The **no** *form of the command disables ICMP rate-limiting on the specified interface(s). (Default:* **Disabled***.)*

1 - 99: *Values in this range allow ICMP traffic as a percentage of the bandwidth available on the interface.*

0 : *This value causes an interface to drop all incoming ICMP traffic , and is not recommended. Refer to the Caution on page 24.*

Note: *ICMP Rate-Limiting is not supported on meshed ports. (Rate-limiting can reduce the efficiency of paths through a mesh domain).*

For example, either of the following commands configures an inbound rate limit of 1% on ports A3 - A5, which are used as network edge ports:

```
HPswitch (config)# int a3-a5 rate-limit icmp 1
HPswitch (eth-A3-A5)# rate-limit icmp 1
```

Displaying the Current Rate-Limit Configuration. This command displays the per-interface rate-limit configuration in the running-config file.

Syntax: show rate-limit icmp [*port-list* | *trunk-list*]

Without **[*port-list* / *trunk-list*]**, *this command lists the ICMP rate-limit configuration for all ports or trunks on the switch. With* **[*port-list* / *trunk-list*]**, *this command lists the rate-limit configuration for the specified interface(s). This command operates the same way in any CLI context.*

For example, if you wanted to view the rate-limiting configuration on the first six ports in the module in slot "B":

```
HPswitch(config)# show rate-limit icmp b1-b6

 Inbound ICMP Rate Limit Maximum Percentage

        | Rate
  Port  | Limit
  ----- + --------
  B1    | Disabled
  B2    | 1
  B3    | 1
  B4    | 1
  B5    | 1
  B6    | Disabled
```

Ports B2-B5 are configured with an ICMP rate limit of 1 %. (Ports B1 and B6 are not configured for ICMP rate-limiting.)

Figure 13. Example of Listing the Rate-Limit Configuration

The **show running** command displays the currently applied setting for any interfaces in the switch configured for ICMP rate limiting. The **show config** command displays this information for the configuration currently stored in the startup-config file. (Note that configuration changes performed with the CLI, but not followed by a **write mem** command do not appear in the startup-config file.)

```
HPswitch(config)# show config status
Running configuration is same as the startup configuration.

HPswitch(config)# show config
Startup configuration:

; J4850A Configuration Editor; Created on release #E.10.00

hostname "HPswitch"
module 2 type J8161A
module 4 type J8161A
snmp-server community "public" Unrestricted
vlan 1
   name "DEFAULT_VLAN"
   untagged B1-B24,D1-D24
   ip address dhcp-bootp
   exit
access-controller vlan-base 2000
interface B2
   rate-limit icmp 1
   exit
interface B3
   rate-limit icmp 1
   exit
interface B4
   rate-limit icmp 1
   exit
interface B5
   rate-limit icmp 1
   exit
```

The **show config status** command compares the content of the startup-config and running-config files and prints a report.

Ports B2-B5 are configured with an ICMP rate limit of 1 %.

Figure 14. Example of ICMP Rate-Limit Settings Listed in the "show running" Output

ICMP Rate-Limiting Trap and Event Log Messages. If the switch detects a volume of inbound ICMP traffic on a port that exceeds the ICMP rate-limit configured for that port, it generates one SNMP trap and one informational Event Log message to notify the system operator of the condition. (The trap and Event Log message are sent within two minutes of the when the event occurred on the port.) For example:

```
I 06/30/05 11:15:42 RateLim: ICMP traffic exceeded configured limit on
port 1
```

These trap and Event Log messages provide an advisory that inbound ICMP traffic on a given interface has exceeded the configured maximum. The additional ICMP traffic is dropped, but the excess condition may indicate an infected host (or other traffic threat or network problem) on that interface. The system operator should investigate the attached devices or network conditions further.

The switch does not send more traps or Event Log messages for excess ICMP traffic on the affected port until the system operator resets the port's ICMP trap function. The reset can be done through SNMP from a network management station or through the CLI with the following **setmib** command.

Syntax: setmib hpIcmpRatelimitPortAlarmflag.< *internal-port-#* > -i 1

On a port configured with ICMP rate-limiting, this command resets the ICMP trap function, which allows the switch to generate a new SNMP trap and an Event Log message if ICMP traffic in excess of the configured limit is detected on the port.

For example, an operator noticing an ICMP rate-limiting trap or Event Log message originating with port A1 on a 5300xl switch would use the following **setmib** command to reset the port to send a new message if the condition occurs again.

```
ProCurve(config)# setmib hpicmpratelimitportalarmflag.1 -i 1
```

Operating Notes for ICMP Rate-Limiting

Note on Testing Rate-Limiting

ICMP rate-limiting is byte-based and is applied to the available bandwidth on an interface. If the total bandwidth requested by all ICMP traffic is less than the available, configured maximum rate, then no ICMP rate-limit can be applied. That is, an interface must be receiving more inbound ICMP traffic than the configured bandwidth limit allows. If the interface is configured with both **rate-limit all** and **rate-limit icmp**, then the ICMP limit can be met or exceeded only if the rate limit for all types of inbound traffic has not already been met or exceeded. Also, to test the ICMP limit it is necessary to generate ICMP traffic that exceeds the configured ICMP rate limit. Using the recommended settings—1% for edge interfaces and 5% maximum for core interfaces—it is easy to generate sufficient traffic. However, if you are testing with higher maximums, it is necessary to ensure that the ICMP traffic

volume exceeds the configured maximum. Note also that testing ICMP rate-limiting where inbound ICMP traffic on a given interface has destinations on multiple outbound interfaces, the test results must be based on the received outbound ICMP average aggregate traffic over time.

ICMP rate-limiting is not reflected in counters monitoring inbound traffic because inbound packets are counted before the ICMP rate-limiting drop action occurs.

- **Interface Support:** ICMP rate-limiting is available on all types of ports and trunks on the switches covered by this guide, and at all port speeds configurable for these devices.
- **Rate-Limiting Not Permitted on Mesh Ports:** Either type of rate-limiting can reduce the efficiency of paths through a mesh domain.
- **Monitoring (Mirroring) ICMP Rate-Limited Interfaces:** If monitoring is configured, packets dropped by ICMP rate-limiting on a monitored interface will still be forwarded to the designated monitor port. (Monitoring shows what traffic is inbound on an interface, and is not affected by "drop" or "forward" decisions.)
- **ICMP Rate-Limit Imposes an Average Bandwidth Limit:** The configured ICMP rate limit on an interface reflects the permitted *average* forwarding rate for ICMP traffic from the interface to the switching fabric. (Note that while occasional bursts of traffic above the configured rate may be observed, the average rate will conform to the configured limit). Rate-Limiting is packet-based, and is calculated internally as the maximum number of 64-byte packets that can be forwarded within the configured bandwidth percentage. Where traffic includes packets larger than 64 bytes, actual average rates may be lower than the configured rate. Also, ICMP rate-limiting reflects the available percentage of an interface's entire inbound bandwidth. The rate of inbound flow for traffic of a given priority and the rate of flow from an ICMP rate-limited interface to a particular queue of an outbound interface are not measures of the actual ICMP rate limit enforced on an interface.
- **Network Stress Conditions:** Under normal network operating conditions, ICMP rate-limiting limits inbound traffic on an interface to no more than the configured level. However, under network stress conditions, the interface may allow occasional, brief bursts of inbound traffic forwarding that exceed the configured rate.
- **Below-Maximum Rates:** ICMP rate-limiting operates on a per-interface basis, regardless of traffic priority. Configuring ICMP rate-limiting on an interface where other features affect inbound port queue behavior (such as flow control) can result in the interface not achieving its configured ICMP rate-limiting maximum. For example, in some situations with flow control configured on an ICMP rate-limited interface, there can be enough "back pressure" to hold high-priority inbound traffic from the upstream device or application to a rate that does not allow bandwidth for lower-priority ICMP traffic. In this case, the inbound traffic flow may not permit the forwarding of ICMP traffic into the switch fabric from the rate-limited interface. (This behavior is termed "head-of-line blocking" and is a well-known problem with flow-control.) In cases where both types of rate-limiting (**rate-limit all** and **rate-limit icmp**) are configured on the same interface, this situation is more likely to occur. In another type of situation, an outbound interface can become oversubscribed by traffic received from multiple ICMP rate-limited inter-

faces. In this case, the actual rate for traffic on the rate-limited interfaces may be lower than configured because the total traffic load requested to the outbound interface exceeds the interface's bandwidth, and thus some requested traffic may be held off on inbound.

- **Optimum Rate-Limiting Operation:** Optimum rate-limiting occurs with 64-byte packet sizes. Traffic with larger packet sizes can result in performance somewhat below the configured inbound bandwidth. This is to ensure the strictest possible rate-limiting of all sizes of packets.
- **Heavy Memory Usage:** Combinations of intensive QoS, rate-limiting, and/or IDM ACL service demands on a switch can impose heavy memory usage on the switch's dynamic hardware rule processor, and can sometimes result in slower system performance. In such cases, moving support for some of the service load to other devices can improve performance.
- **Outbound Traffic Flow:** Configuring ICMP rate-limiting on an interface does not control the rate of outbound traffic flow on the interface.
- **Traffic Filters on Rate-Limited Interfaces:** Configuring a traffic filter on an interface does not prevent the switch from including filtered traffic in the bandwidth-use measurement for either type of rate-limiting (ICMP or all). That is, where rate-limiting and traffic filtering are configured on the same interface, the inbound, filtered traffic is included in the bandwidth measurement for calculating when the limit has been reached. Traffic filters include:
  - ACLs
  - Source-Port filters
  - Protocol filters
  - Multicast filters
- **Determining the 5300xl Switch Port Number Used in ICMP Port Reset Commands To Enable Excess ICMP Traffic Notification Traps and Event Log Messages:** Use the internal port numbers described in this section with the **setmib** command described on page 29. The port number included in the command corresponds to the internal number the switch maintains for the designated port, and not the port's external (slot/number) identity. To match the port's external slot/number to the internal port number, use the **walkmib ifDescr** command, as shown in the following figure:

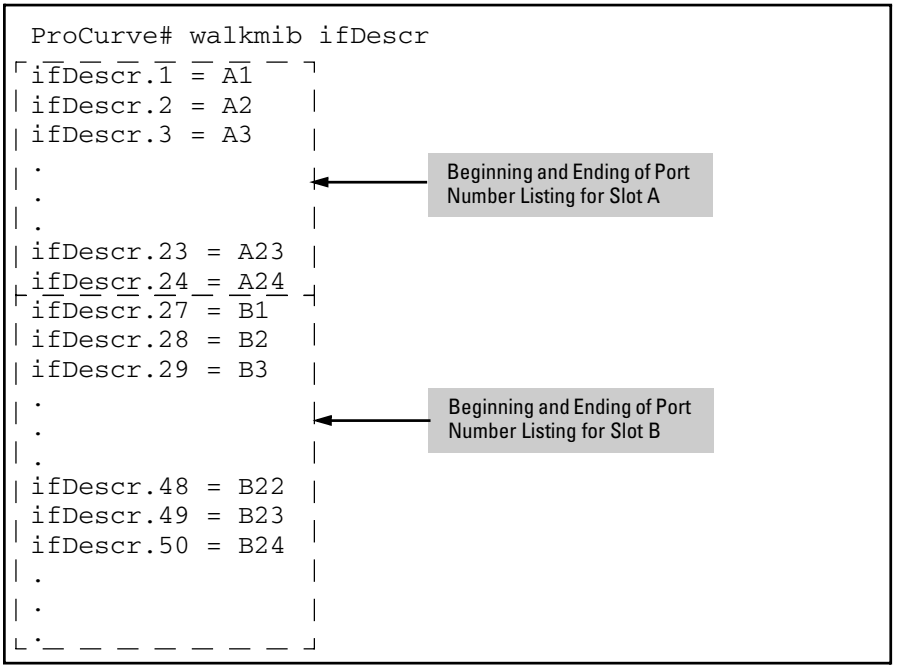

Figure 15. Matching Internal Port Numbers to External Slot/Port Numbers

RADIUS Accounting for Switch Access Through Web Authentication and MAC Authentication Sessions

Prior to software release E.10.02, the switch supported RADIUS accounting for 802.1X authenticated client sessions. Beginning with release E.10.02, the switch also supports RADIUS accounting for Web Authentication and MAC Authentication sessions. To configure RADIUS accounting on the switch, refer to the chapter titled "RADIUS Authentication and Accounting" in the Access Security Guide for your 5300xl switch.

RADIUS-Assigned Access Control Lists

This feature uses RADIUS-assigned, per-port ACLs for Layer-3 filtering of inbound IP traffic from authenticated clients. A given RADIUS-assigned ACL is identified by a unique username/password pair or client MAC address, and applies only to traffic from clients that authenticate with the same unique credentials. ACL services for an authenticated client include filtering inbound IP traffic based on destination and/or IP traffic type (such as TCP and UDP traffic) and traffic counter options. Implementing the feature requires:

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

A RADIUS-assigned ACL is a type of extended ACL that filters IP traffic inbound on a port from any source (and, optionally, of any specific IP application or protocol type) to a single destination IP address, a group of contiguous IP addresses, an IP subnet, or any IP destination.

This feature is designed to accept dynamic configuration of a RADIUS-based ACL on an individual port on the network edge to filter traffic from an authenticated end-node client. Using RADIUS to apply per-port ACLs to edge ports enables the switch to filter IP traffic coming from outside the network, thus removing unwanted traffic as soon as possible and helping to improve system performance. Also, applying RADIUS-assigned ACLs to ports on the network edge is likely to be less complex than using VLAN-based ACLs in the network core to filter unwanted traffic that could have been filtered at the edge.

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

Note

A RADIUS-assigned ACL filters all inbound IP traffic from an authenticated client on a port, regardless of whether the traffic is to be switched or routed. (VLAN-based ACLs configurable on 5300xl switches filter only routed traffic and traffic with a destination address—DA—on the switch itself.)

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

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

Table 2, below, highlights several key differences between the static ACLs configurable on 5300xl switch VLANs and the dynamic ACLs that can be assigned to individual ports by a RADIUS server.

Table 2. Contrasting Dynamic and Static ACLs

RADIUS-Based (Dynamic) ACLs	VLAN-Based (Static) ACLs
Configured in client accounts on a RADIUS server.	Configured in the switch itself.
Designed for use on the edge of the network where filtering of inbound traffic is most important and where clients with differing access requirements are likely to use the same port at different times.	Designed for general use where the filtering needs for traffic to or from connected devices are predictable and largely static.
Implementation requires client authentication.	Client authentication not a factor.
Identified by the credentials (username/password pair or the MAC address) of the specific client the ACL is intended to service.	Identified by a number in the range of 1-199 or an alphanumeric name.
Supports dynamic assignment to filter only the inbound IP traffic from an authenticated client on the port to which the client is connected. (Traffic can be routed or switched, and includes traffic having a DA on the switch itself.)	Supports static assignments to filter either inbound or outbound for all ports in the assigned VLAN, routed IP traffic, and inbound IP traffic having a DA on the switch itself.
When the authenticated client session ends, the switch removes the RADIUS-assigned ACL from the client port.	Remains statically assigned to the VLAN unless removed by a **no vlan < *vid* > ip access-group** CLI command.
Supports a maximum of two RADIUS-based ACLs on a port. (Each ACL supports one authenticated client.)	Supports one inbound ACL and one outbound ACL per-VLAN.
Supports only extended ACLs. (Refer to Terminology.)	Supports standard, extended, and connection-rate ACLs, and applies these ACLs to traffic on all ports belonging to the VLAN..
The ACL filters only the IP traffic it receives inbound from the authenticated client corresponding to that ACL, and does not filter traffic inbound from other authenticated clients.(The traffic source is not a configurable setting.)	An ACL applied inbound on a VLAN filters all IP traffic received on any member port from any source in the same VLAN, as long as the traffic is either routed by the switch to another VLAN or subnet, or has a DA on the switch itself. An ACL applied outbound on a VLAN filters all routed IP traffic leaving the switch on any member port.
Can contain up to 30 ACEs.	Can contain up to 1024 ACEs per 5300xl switch.
Requires client authentication by a RADIUS server configured to dynamically assign an ACL to the client port, based on client credentials.	Configured in the switch and statically applied to filter IP traffic on all ports in the specified VLAN, regardless of other factors.
ACEs allow a counter (**cnt**) option that causes a counter to increment when there is a packet match.	ACEs allow a **log** option that generates a log message whenever there is a packet match with a "deny" ACE.

Terminology

ACE: See Access Control Entry, below.

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

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

ACL: See Access Control List, below.

Access Control List (ACL): A list (or set) consisting of one or more explicitly configured Access Control Entries (ACEs) and terminating with an implicit "deny" default which drops any packets that do not have a match with any explicit ACE in the named ACL.

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

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

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

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

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

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

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

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

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

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

Wildcard: The part of a mask that indicates the bits in a packet's IP addressing that do not need to match the corresponding bits specified in an ACL. See also **ACL Mask** on page 35.

Caution Regarding the Use of Source Routing

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

General Operation

An ACL is a list of one or more Access Control Entries (ACEs), where each ACE consists of a matching criteria and an action (permit or deny). These ACEs are designed to control the network access privileges of an authenticated client. A RADIUS-based ACL applies only to the inbound traffic from the client whose authentication triggers the ACL assignment to the client port.

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

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

Notes

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

On a given port, RADIUS-based ACL filtering occurs only for the inbound traffic from the client whose authentication configuration on the server includes a RADIUS-based ACL. Inbound traffic from another authenticated client (on the same port) whose authentication configuration on the server does not include a RADIUS-based ACL will not be filtered by a RADIUS-based ACL assigned to the port for any other authenticated client.

The Packet-filtering Process

Sequential Comparison and Action. When an ACL filters a packet, it sequentially compares each ACE's filtering criteria to the corresponding data in the packet until it finds a match. The action indicated by the matching ACE (deny or permit) is then performed on the packet.

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

Example. Suppose the ACL in figure 16 is assigned to filter the traffic from an authenticated client on a given port in the switch:

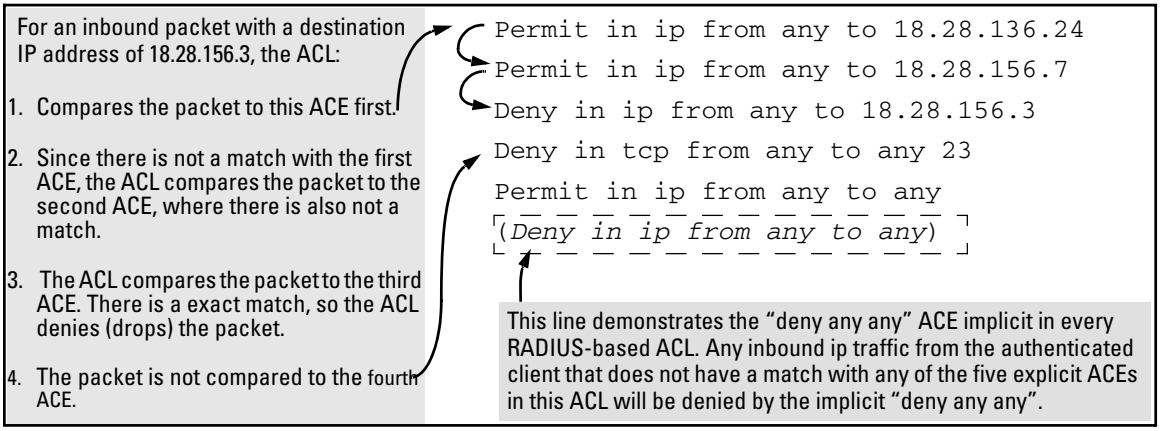

Figure 16. Example of Sequential Comparison

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

Note

If a RADIUS-based ACL permits an authenticated client's inbound IP packet, but the client port belongs to a VLAN for which there is an inbound, VLAN-based ACL configured on the switch, then the packet will also be filtered by the VLAN-based ACL.

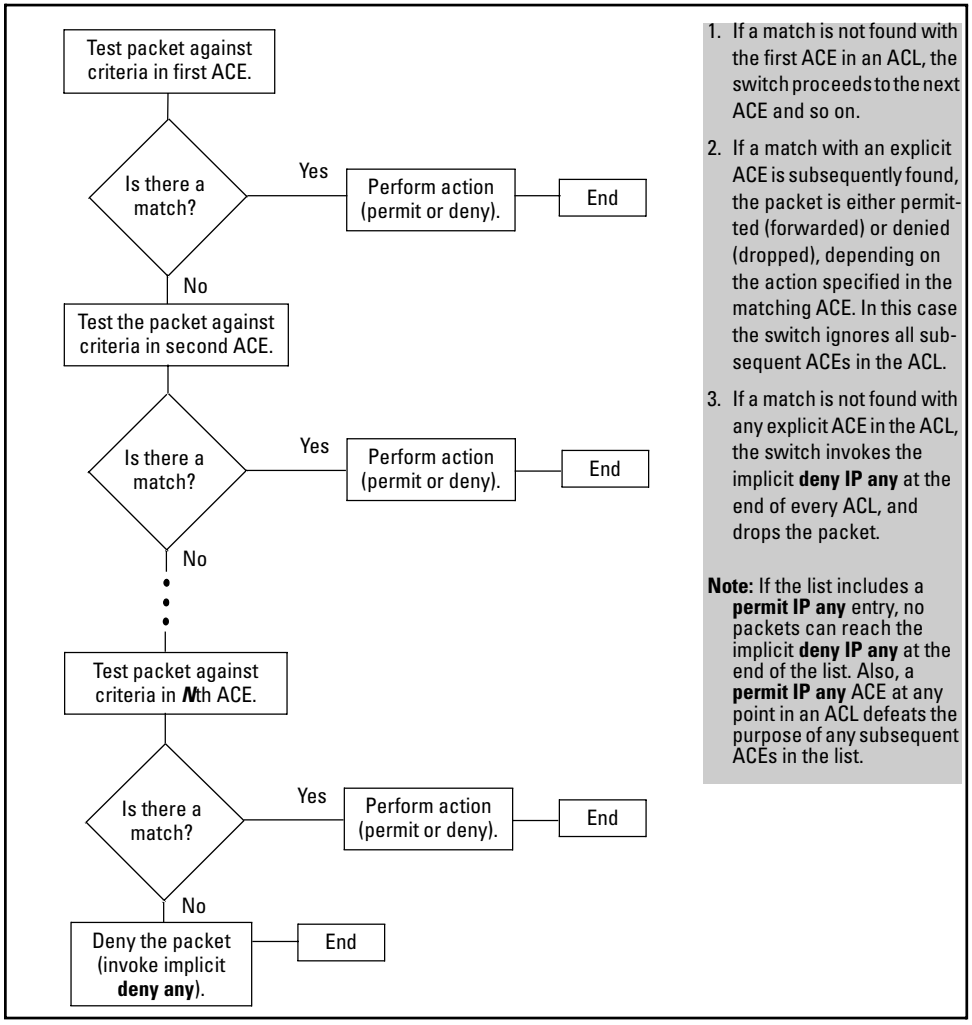

Figure 17. The Packet-Filtering Process in an ACL with *N* Entries (ACEs)

Note

The order in which an ACE occurs in an ACL is significant. For example, if an ACL contains six ACEs, but the first ACE is a "permit IP any", then the ACL permits all IP traffic, and the remaining ACEs in the list do not apply, even if they specify criteria that would make a match with any of the traffic permitted by the first ACE.

For example, suppose you want to configure a RADIUS-based ACL to invoke these policies in the 11.11.11.0 network:

1. Permit inbound client traffic with a DA of 11.11.11.42.
2. Permit inbound Telnet traffic for DA 11.11.11.101.
3. Deny inbound Telnet traffic for all other IP addresses in the 11.11.11.0 network.
4. Permit inbound HTTP traffic for any IP address in the 11.11.11.0 network.
5. Deny all other inbound traffic.

The following ACL model, when invoked by a client authenticating with the credentials configured in the RADIUS server for this ACL, supports the above case:

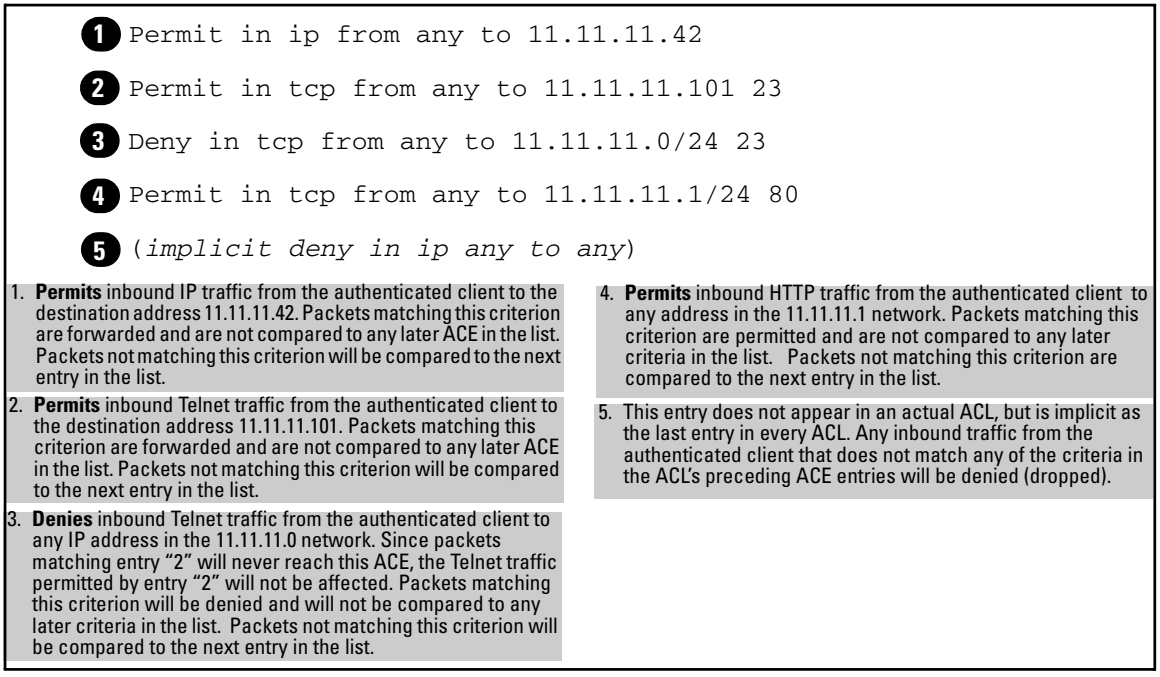

Figure 18. Example of How a RADIUS-Based ACL Filters Packets

It is important to remember that RADIUS-based ACLs include an implicit "deny IP any any". That is, packets received inbound from an authenticated client that the ACL does not *explicitly* permit or deny will be *implicitly* denied, and therefore dropped instead of forwarded. If you want the port to permit all inbound IP traffic (from the authenticated client) that the ACL does not explicitly permit or deny, insert a **permit in ip from any to any** ("permit any any") as the last explicit entry in the ACL.

Overriding the Implicit "deny IP any any". If you want an ACL to permit any routed packets that are not explicitly denied by other entries in the ACL, you can do so by configuring a **permit any** entry as the last entry in the ACL. Doing so permits any packet not explicitly denied by earlier entries.

General Steps

These steps suggest a process for using ACLs to establish client access policies. The topics following this section provide details.

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

Determining Traffic Policies

This section assumes that the RADIUS server needed by a client for authentication and ACL assignments is accessible from any switch that authorized clients may use.

Begin by defining the policies you want an ACL to enforce for a given client or group of clients. This includes the type of IP traffic permitted or not permitted from the client(s) and the areas of the network the client(s) are authorized or not authorized to use.

- What traffic should you permit for the client or group? In some cases you will need to explicitly identify permitted traffic. In other cases, depending on your policies, you can insert a **permit any/any** entry at the end of the ACL so that all IP traffic not specifically matched by earlier entries in the list will be permitted. This may be the best choice for an ACL that begins by defining the inbound client IP traffic that should be dropped.
- What traffic must be explicitly blocked for the client or group? This can include requests to access to "off-limits" subnets, unauthorized access to the internet, access to sensitive data storage or restricted equipment, and preventing the use of specific TCP or UDP applications such as Telnet, SSH, and web browser access to the switch.

- What traffic can be blocked simply by relying on the implicit **deny any/any** that is automatically included at the end of every ACL? This can reduce the number of entries needed in an ACL.
- Is it important to keep track of the number of matches for a particular client or ACE? If so, you can use the optional **cnt** (counter) feature in ACEs where you want to know this information. This is especially useful if you want to verify that the switch is denying unwanted client packets. (Note that configuring a high number of counters can exhaust the counter resources. Refer to table 3 on page 44.)

Caution

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

Planning the ACLs Needed To Enforce Designated Traffic Policies

This section can help in understanding how to order the ACEs in a RADIUS-based ACL and in understanding how clients and the switch operate in this dynamic environment.

Guidelines for Structuring a RADIUS-Based ACL.

- The sequence of ACEs is significant. When the switch uses an ACL to determine whether to permit or deny a packet on a particular VLAN, it compares the packet to the criteria specified in the individual Access Control Entries (ACEs) in the ACL, beginning with the first ACE in the list and proceeding sequentially until a match is found. When a match is found, the switch applies the indicated action (permit or deny) to the packet. This is significant because, when a match is found for a packet, subsequent ACEs in the same ACL will not be used for that packet, regardless of whether they match the packet.
- **Inbound Traffic Only:** RADIUS-based ACLs filter only the inbound IP traffic from an authenticated client for which an ACL has been configured on the appropriate RADIUS server.
- **Result of an ACE/Packet Match:** The first match of a given packet to an ACE dictates the action for that packet. Any subsequent match possibilities are ignored.
- **Explicitly Permitting Any IP Traffic:** Entering a **permit in ip from any to any** (permit any any) ACE in an ACL permits all IP traffic not previously permitted or denied by that ACL. Any ACEs listed after that point do not have any effect.

- **Explicitly Denying Any IP Traffic:** Entering a **deny in ip from any to any** ACE in an ACL denies all IP traffic not previously permitted or denied by that ACL. Any ACEs listed after that point have no effect.
- **Implicitly Denying Any IP Traffic:** For any packet being filtered by an ACL, there will always be a match. Included in every ACL is an implicit **deny in ip from any to any**. This means that the ACL denies any IP packet it filters that does not have a match with an explicitly configured ACE. Thus, if you want an ACL to permit any packets that are not explicitly denied, you must configure **permit in ip from any to any** as the last explicit ACE in the ACL. Because, for a given packet, the switch sequentially applies the ACEs in an ACL until it finds a match, any packet that reaches the **permit in ip from any to any** entry will be permitted, and will not reach the implicit **deny in ip from any to any** ACE that is included at the end of the ACL. For an example, refer to figure 18 on page 40.
- Determine the order in which you want the individual ACEs in the ACL to filter inbound traffic from a client. A general guideline is to arrange the ACEs in the expected order of decreasing application frequency. This will result in the most prevalent traffic types finding a match earlier in the ACL than traffic types that are more infrequent, thus saving processing cycles.

Operating Rules for RADIUS-Based ACLs

- **Relating a Client to a RADIUS-Based ACL:** A RADIUS-based ACL for a particular client must be configured in the RADIUS server under the authentication credentials the server should expect for that client. (If the client must authenticate using 802.1X and/or Web Authentication, the username/password pair forms the credential set. If authentication is through MAC Authentication, then the client MAC address forms the credential set.) For more on this topic, refer to "Configuring an ACL in a RADIUS Server" on page 45.
- **Multiple Clients Using the Same Username/Password Pair:** Multiple clients using the same username\password pair will use duplicate instances of the same ACL.
- **Limits for RADIUS-Based ACLs, Associated ACEs, and Counters:**

  Table 3 describes limits the switch supports in ACLs applied by a RADIUS server. Exceeding a limit causes the related client authentication to fail.

Table 3. Limits Affecting RADIUS-Based ACL Applications

Item	Limit	Notes
Maximum Number of Authenticated Sessions Per-Port Using RADIUS-based ACLs	2	A port supports a maximum of two ACLs (or two instances of the same ACL) on a given port at the same time. *This rule does not affect the number of authenticated clients a port supports (32); only the number of authenticated clients using RADIUS-based ACLs.* If two authenticated clients are already using RADIUS-based ACLs on a port and a third client on the same port attempts to authenticate with a RADIUS server account that includes an ACL assignment, the attempt will fail.
Maximum Number of (internal) ACEs Per-Port, and Maximum Number of (internal) ACEs Per-ACL	30	Depending on how a RADIUS-assigned ACE is formed, it can consume multiple internal ACEs. A RADIUS-assigned ACE that does not specify TCP or UDP port numbers uses one internal ACE. However, an ACE that includes TCP or UDP port numbers uses one or more internal ACE resources, depending on the port number groupings. A single TCP or UDP port number or a series of contiguous port numbers comprise one group. For example, "80" and "137-146" each form one group. "135, 137-140, 143" in a given ACE form three groups. The following ACE examples illustrate how the switch applies internal ACE usage. **Examples of Single and Multiple (Internal) ACEs Per-Port** — **Internal ACEs** deny in ip from any to any — 1 deny in tcp from any to any — 1 deny in tcp from any to any 80 — 1 permit in tcp from any to any 135, 137-146, 445 — 3 permit in tcp from any to any 135-137, 139, 141, 143, 146, 445 — 6 permit in tcp from any to any 135-146, 445 — 2 Where two authenticated clients are using RADIUS-based ACLs on the same port, the total number of ACEs in both active sessions cannot exceed the maximum.
Maximum Number of Characters in a single ACE	80	—
Maximum Number of (optional) Internal Counters Used Per-Module	100	Depending on how an ACE is formed, using the **cnt** (counter) option consumes one or more internal counters. Using a counter in an ACE that does not specify TCP or UDP port numbers uses one counter. Using a counter in an ACE that includes TCP or UDP port numbers uses one or more counters, depending on the port number groupings. A single TCP or UDP port number or a series of contiguous port numbers comprise one group. For example, "80" and "137-146" each form one group. "135, 137-140, 143" in a given ACE form three groups. The following ACE examples illustrate how the switch calculates internal counter groups. **Examples of ACEs Employing Counters** — **Internal Counters** deny in ip from any to any cnt — 1 deny in tcp from any to any cnt — 1 deny in tcp from any to any 80 cnt — 1 permit in tcp from any to any 135, 137-146, 445 cnt — 3 permit in tcp from any to any 135-137, 139, 141, 143, 146, 445 cnt — 6 permit in tcp from any to any 135-146, 445 cnt — 2

- **Effect of VLAN-Based ACLs Configured on the Switch:** A port receiving a dynamic, RADIUS-based ACL assignment can also belong to a VLAN for which there is an inbound ACL statically configured (on the switch). In this case, an IP packet permitted by the RADIUS-based ACL will also be filtered by the VLAN-based ACL if the inbound client packets are routed or have a DA on the switch itself. If the RADIUS-based ACL permits the packet, but the VLAN-based, inbound ACL denies the packet, then the packet is dropped. If the RADIUS-based ACL denies the packet, then the packet is dropped and does not reach the VLAN-based, inbound ACL. (RADIUS-based ACLs operate only on inbound IP traffic, and are not a factor for the traffic filtered by VLAN-based, outbound ACLs.)

- **A RADIUS-Based ACL Affects Only the Inbound Traffic from a Specific, Authenticated Client:** A RADIUS-based ACL assigned to a port as the result of a client authenticating on that port applies only to the inbound traffic received on that port from that client. It does not affect the traffic received from any other authenticated clients on that port, and does not affect any outbound traffic on that port.

Configuring an ACL in a RADIUS Server

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

Elements in a RADIUS-Based ACL Configuration. A RADIUS-based ACL configuration in a RADIUS server has the following elements:

- vendor and ACL identifiers:
  - ProCurve (HP) Vendor-Specific ID: 11
  - Vendor-Specific Attribute for ACLs: 61 (string = HP-IP-FILTER-RAW)
  - Setting: HP-IP-FILTER-RAW = < “permit” or “deny” ACE >

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

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

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

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

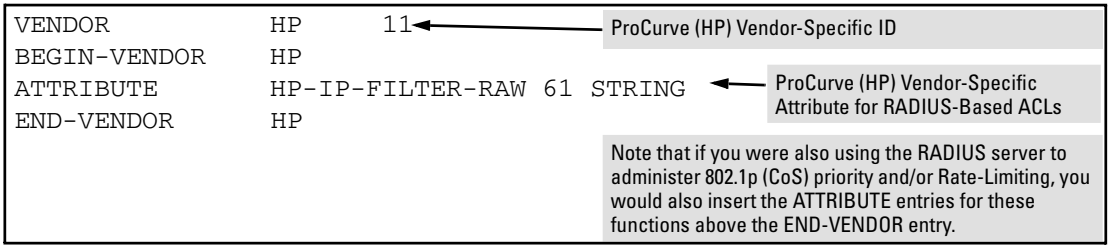

Figure 19. Example of Configuring the VSA for RADIUS-Based ACLs in a FreeRADIUS Server

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

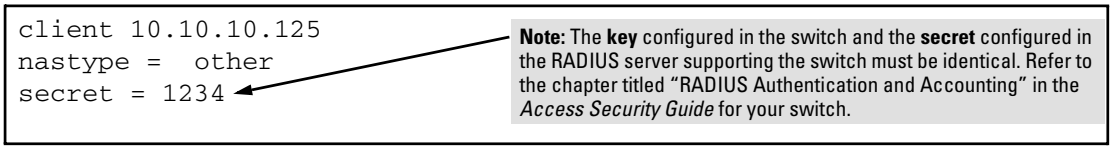

Figure 20. Example of Configuring the Switch's Identity Information in a FreeRADIUS Server

3. For a given client username/password pair or MAC address, create an ACL by entering one or more ACEs in the FreeRADIUS "users" file. Enter the ACEs in an order that promotes optimum traffic management and conservation of system resources, and remember that every ACL you create automatically includes an implicit **deny in ip from any to any** ACE. (Refer to "Guidelines for Structuring a RADIUS-Based ACL" on page 42.) For example, suppose that you wanted to create identical ACL support for the following:

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

   The ACL in this example must achieve the following:

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

   To configure the above ACL, you would enter the username/password and ACE information shown in figure 21 into the FreeRADIUS **users** file.

Note

For syntax details on RADIUS-based ACLs, refer to "Format Details for ACEs Configured in a RADIUS-Based ACL" on page 47.

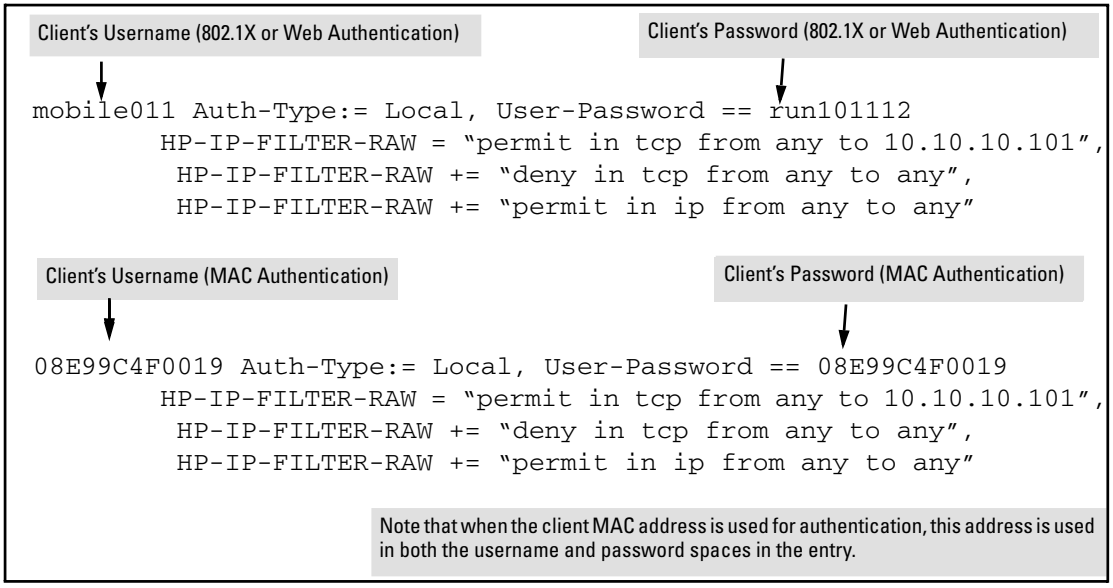

Figure 21. Example of Configuring the FreeRADIUS Server To Support ACLs for the Indicated Clients

Format Details for ACEs Configured in a RADIUS-Based ACL.

Any instance of a RADIUS-Based ACL is structured to filter authenticated client traffic as follows:

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

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

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

< permit | deny >: *Specifies whether to forward or drop the identified IP traffic type from the authenticated client.*

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

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

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

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

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

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

to : *Required destination keyword.*

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

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

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

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

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

[cnt]: *Optional counter specifier for a RADIUS-based ACL. When used in an ACL, the counter increments each time there is a "match" with a permit or deny ACE. This option requires that you configure the switch for RADIUS accounting. (Refer to the entry describing the maximum number of (optional) internal counters in the table on page 44.)*

Configuring the Switch To Support RADIUS-Based ACLs

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

1. Configure RADIUS operation on the switch:

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

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

2. Configure RADIUS network accounting on the switch (optional). RADIUS network accounting is necessary to retrieve counter information if the **cnt** (counter) option (described on page 48) is included in any of the ACEs configured on the RADIUS server.

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

   For more on RADIUS accounting, refer to *the chapter titled "RADIUS Authentication and Accounting" in the* Access Security Guide for your switch.

Note

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

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

   802.1X Option:

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

   These commands configure 802.1X port-based access control on the switch, and activates this feature on the specified ports. For more on 802.1X configuration and operation, refer to the chapter titled "Configuring Port-Based and Client-Based Access Control" in the *Access Security Guide* for your switch.

   MAC Authentication Option:

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

   This command configures MAC authentication on the switch and activates this feature on the specified ports. For more on MAC authentication, refer to the chapter titled "Web and MAC Authentication" in the *Access Security Guide* for your switch.

   Web Authentication Option:

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

   This command configures Web authentication on the switch and activates this feature on the specified ports. For more on Web authentication, refer to the chapter titled "Web and MAC Authentication" in the *Access Security Guide* for your switch.

Displaying the Current RADIUS-Based ACL Activity on the Switch

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

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

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

Note: *If there are no ACLs currently assigned to any port in* < *port-list* >*, executing this command returns only the system prompt. If a client authenticates but the server does not return a RADIUS-based ACL to the client port, then the server does not have a valid ACL configured and assigned to that client's authentication credentials.*

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

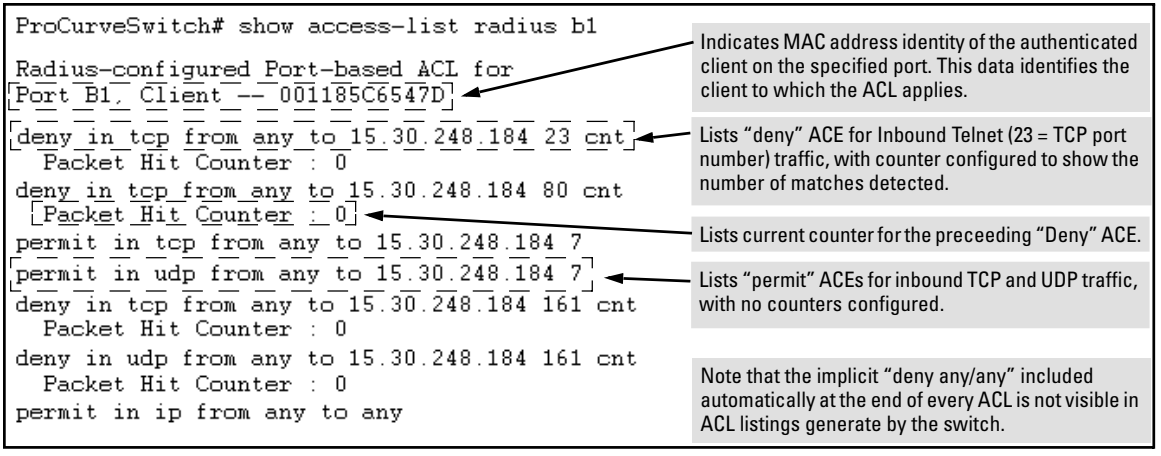

Figure 22. Example Showing a RADIUS-Based ACL Application to a Currently Active Client Session

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

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

Port: *Port number of port configured for authentication.*

Status: *Port connection status:*

Open = *active connection with an external device*

Closed = *no active connection with an external device*

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

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

No-override: *Indicates that no RADIUS-assigned 802.1p priority is currently active on the indicated port. (For more on traffic prioritization for the 5300xl switches, refer to the chapter titled "Quality of Service (QoS): Managing Bandwidth More Effectively" in the Advanced Traffic Management Guide for your switch.)*

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

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

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

0 - 100: *Indicates that the displayed rate-limit has been assigned by a RADIUS server to inbound traffic on the indicated port for a currently active, authenticated client session. This assignment remains active until the session ends.*

RADIUS ACL Applied?: *Indicates whether a RADIUS-assigned ACL is currently active on the port.*

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

No: *There is no RADIUS-assigned ACL currently active on the indicated port.*

```
HPswitch# show port-access authenticator b1

 Port Access Authenticator Status

  Port-access authenticator activated [No] : Yes

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

Indicates a RADIUS ACL is currently applied as part of an active session with an authenticated client.

Figure 23. Example of Output Showing Current RADIUS-Applied Features

Event Log Messages

Message	Meaning
`ACE parsing error, permit/deny keyword <ace-#> client <mac-address> port <port-#>.`	Notifies of a problem with the `permit/deny` keyword in the indicated ACE included in the access list for the indicated client on the indicated switch port.
`Could not add ACL entry.`	Notifies that the ACE entry could not be added to the internal ACL storage.
`Could not create ACL entry.`	Notifies that the ACL could not be added to the internal ACL storage.
`Could not add ACL, client mac <mac-address>port <port-#>, at max per-port ACL quantity.`	Notifies that the ACL could not be added because the per-port ACL quantity would be exceeded.
`ACE parsing error, IN keyword, <ace-#> client <mac-address> port <port-#>.`	Notifies of a problem with the `IN` keyword in the indicated ACE of the access list for the indicated client on the indicated switch port.
`ACE parsing error, protocol field, <ace-#> client <mac-address> port <port-#>.`	Notifies of a problem with the protocol field in the indicated ACE of the access list for the indicated client on the indicated switch port.
`ACE parsing error, FROM keyword, <ace-#> client <mac-address> port <port-#>.`	Notifies of a problem with the `FROM` keyword in the indicated ACE of the access list for the indicated client on the indicated switch port.
`ACE parsing error, ANY keyword, <ace-#> client <mac-address> port <port-#>.`	Notifies of a problem with the `ANY` keyword in the indicated ACE of the access list for the indicated client on the indicated switch port.
`ACE parsing error, TO keyword, <ace-#> client <mac-address> port <port-#>.`	Notifies of a problem with the `TO` keyword in the indicated ACE of the access list for the indicated client on the indicated switch port.

Message	Meaning
`ACE parsing error, destination IP, <ace-#> client <mac-address> port <port-#>.`	Notifies of a problem with the destination IP field in the indicated ACE of the access list for the indicated client on the indicated switch port.
`ACE parsing error, tcp/udp ports, <ace-#> client <mac-address> port <port-#>.`	Notifies of a problem with the TCP/UDP port field in the indicated ACE of the access list for the indicated client on the indicated switch port.
`Rule limit per ACL exceeded. <ace-#> client <mac-address> port <port-#>.`	Notifies that an ACL has too many rules. A maximum of 30 (internal) ACEs are allowed per ACL. Refer to table 3 on page 44.
`Duplicate mac. An ACl exists for client. Deauthenticating second. client <mac-address> port <port-#>.`	Notifies that an ACL for this mac on this port already exists.
`Invalid Access-list entry length, client <mac-address> port <port-#>.`	Notifies that the string configured for an ACE entry on the Radius server exceeds 80 characters.
`Memory allocation failure for IDM ACL.`	Notifies of a memory allocation failure for a RADIUS-based ACL. **User Action?**
`ACE limit per port exceeded. client <mac-address> port <port-#>.`	Notifies that the maximum number of ACEs (30) allowed on the port was exceeded.
`Exceeded counter per slot limit. client <mac-address> port <port-#>.`	Notifies that the internal counter (cnt) limit of 100 per module was exceeded on port <*port-#*>. Refer to table 3 on page 44.

Causes of Client Deauthentication Immediately After Authenticating

- ACE formatted incorrectly in the RADIUS server
  - "from", "any", or "to" keyword missing
  - An IP protocol number in the ACE exceeds 255.
  - An optional UDP or TCP port number is invalid.
- A RADIUS-Based ACL limit has been exceeded. (Refer to table 3, "Limits Affecting RADIUS-Based ACL Applications" on page 44.)
  - The allowed maximum of two RADIUS-assigned ACLs has already been reached on the port through which the deauthenticated client is trying to access the network. (Each client requiring a RADIUS-assigned ACL is a separate instance, even if multiple clients are assigned the same ACL.)
  - For a given port on a given module, the latest client authentication includes a RADIUS-Based ACL assignment exceeding the maximum number of ACEs allowed on the module.
  - An ACE in the ACL for a given authenticated client exceeds 80 characters.
  - An ACL assigned to an authenticated client causes the number of optional counters needed on the module supporting the client's port to exceed the per-module maximum (100).

Release E.09.22 Enhancements

The table below summarizes the enhancements made in this release. New features are described in detail in the sections following the table.

Enhancement	Overview
Supports ProCurve Gigabit 1000Base-T Mini-GBIC (J8177B)	Pluggable Gigabit transceiver (RJ-45) for up to 100m over Category 5 cable or better.
XRRP Infinite Fail-Back	See "XRRP Infinite Fail-Back for the 5300xl Switches" on page 55 for details.
DHCP Option 82	See "DHCP Option 82" on page 61 for details.

XRRP Infinite Fail-Back for the 5300xl Switches

Introduction

XRRP infinite failback is an optional enhancement to the Series 5300xl XRRP routing feature, and is designed to reduce network disruption due to peer router fail-backs occurring automatically. This is accomplished by configuring XRRP with the infinite fail-back option and then using a manual fail-back command that can be executed at the discretion of a system operator.

Series 5300xl switches using XRRP are considered to be routers, and the term "router" is used to identify them.

Before using this section you should have an understanding of XRRP operation and terminology as described in the chapter titled "Router Redundancy Using XRRP" in the Advanced Traffic Management Guide for your 5300xl device.

Terminology

Fail-Back Router: In a given protection domain, this is the XRRP-enabled router that takes over the routing functions transferred from its XRRP peer in the domain when the peer loses access to one or more of its XRRP VLANs. The fail-back router must have access to all of its XRRP VLANs at the time of the fail-over. See also **Fail-Over Router**.

Fail-Over Router: In a given protection domain, this is an XRRP-enabled router that loses access to one or more of its XRRP VLANs, causing a fail-over of its routing functions to the other XRRP router (peer) in the domain. The peer must have access to all of its XRRP VLANs at the time of the fail-over, and is designated as the "fail-back" router. See also **Fail-Back Router**.

Infinite Fail-Back: The operating mode of an XRRP router in which the router does not automatically allow fail-back when its peer in the protection domain recovers from a fail-over condition.

Permanent Control: The mode associated with infinite fail-back in which a fail-over has occurred in a protection domain and the resulting fail-back router has both primary control and secondary control of the XRRP VLANs in the domain. See also **Primary Control** and **Secondary Control**. (For events that terminate a permanent control mode on an XRRP router, refer to "Router Operation in the Infinite Fail-Back Mode" on page 58.)

Primary Control: The mode in which the XRRP VLANs configured on a router in a protection domain are controlled by that router.

Secondary Control: The mode in which the XRRP VLANs configured on one router in a protection domain are controlled by the other (fail-back) XXRP router in the domain. A fail-back router advertising XRRP packets for the failed (peer) router's backed-up IP addresses has both primary and secondary control of the XRRP VLANs in the domain.

Overview of Infinite Fail-Back Operation

If a fail-over event occurs in an XRRP protection domain, the peer for the failed router automatically takes over the routing function for the failed router. This peer router, which is already the master for its own XRRP VLAN(s) and XRRP MAC address (primary address control), also becomes the master for the XRRP VLAN(s) and XRRP MAC address for the failed router (secondary address control) .

The Problem. Prior to software release E.09.05, if a failed XRRP router recovers access to all of its XRRP VLANs, then a fail-back automatically occurs, which removes secondary address control from the fail-back router and restores control of these addresses to the recovered router. (The fail-back router ceases to advertise XRRP packets for the failed router's backed-up IP addresses.) This automatic fail-back can cause a network slowdown due to a disruption of the TCP connections during the fail-back. In cases where a fail-over/fail-back cycle occurs repeatedly, frequent network disruptions can occur.

The Solution. Beginning with software release E.09.05, you can optionally configure *XRRP infinite fail-back*, which blocks automatic fail-back as long as the fail-back router continues XRRP operation with at least one of its XRRP VLANs remaining up. In this mode, the fail-back router maintains "permanent" primary and secondary address control. This means that recovery of the fail-over router does not automatically result in a fail-back from its peer, and can only occur when either a system operator uses the CLI to force fail-back or there is a system change affecting the fail-back router. (Events that can cause a fail-back to occur are described under "Router Operation in the Infinite Fail-Back Mode" on page 58.)

Causes of Fail-Over and Fail-Back

Fail-Over. An XRRP router fail-over to its peer occurs in these instances:

- The router loses connectivity on *any* XRRP VLAN (and the peer router has maintained connectivity with *all* of its XRRP VLANs).
- None of an XRRP router's multicast advertisements are detected by the peer router within three advertisement intervals (and the peer router has maintained connectivity with at least one of its XRRP VLANs).

Note

If a peer router is unable to support a fail-over due to its own failures, the fail-over does not occur and routing will cease in at least some areas of the network.

Fail-Back. Without infinite fail-back enabled, an XRRP peer router with both primary and secondary control in a protection domain automatically initiates a fail-back when it detects XRRP advertisements indicating that the failed router has recovered access to all of its XRRP VLANs.

With infinite fail-back enabled in the protection domain, a change in the state of the failed router does not initiate a fail-back.

Fail-Over Operation with Infinite Fail-Back Enabled

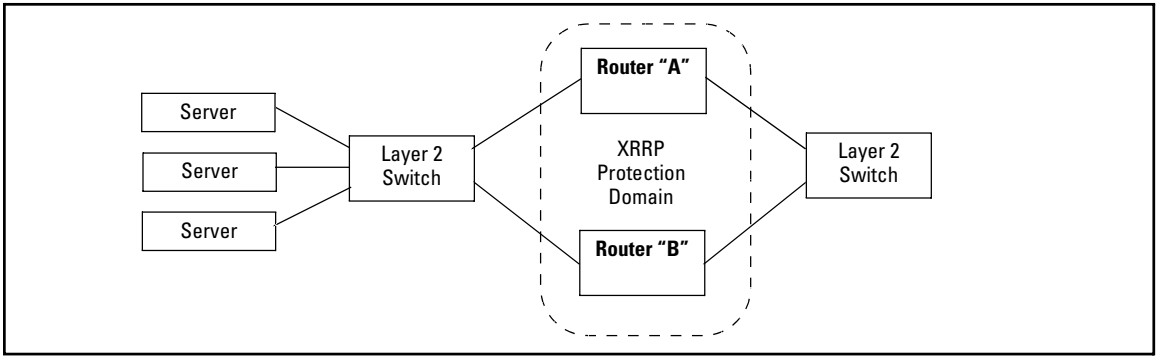

Figure 24. Example of XRRP Routers in a Protection Domain

When XRRP is enabled on router "A" with infinite fail-back already enabled in the configuration, the router immediately waits for a period equal to three times the XRRP advertisement interval to determine the current fail-back mode of its XRRP peer router "B". Depending on the result of the waiting period, router "A" assumes one of the following modes:

- **primary control only:** occurs if router "B" is doing one of the following:
  - running with only primary control mode.
  - running with primary and secondary control mode, but *has not* detected a prior fail-over on router "A". (This can occur if a failed router is replaced with another XRRP router.)
- **primary and secondary control:** occurs if router "B" is in a fail-over mode.
- **fail-over:** occurs if router "B" *has* detected a prior fail-over on router "A" and is in the primary and secondary control mode as a result.

Router Operation in the Fail-Over Mode

An XRRP router that has failed-over to a peer (fail-back router) in a protection domain (where both routers are configured with infinite fail-back), and then recovers, remains in the fail-over mode unless the status of the fail-back router changes as described under "Router Operation in the Infinite Fail-Back Mode", below. This means that changes solely in the status of the failed-over router, such as recovery of all its XRRP VLANs, cannot initiate a fail-back from the peer router.

Router Operation in the Infinite Fail-Back Mode

An XRRP router with permanent (primary and secondary) address control and infinite fail-back enabled will not surrender permanent control (fail-back) to a recovered fail-over router except as described below.

- XRRP traffic is moving between the two routers and a system operator initiates fail-back by using the **xrrp ctrl-transfer** CLI command in either of the following cases:
  - on the fail-back router to force a fail-back to the fail-over router
  - on the fail-over router, provided that traffic can move between the fail-back router and the fail-over router

  (This command allows the fail-over router to resume primary control of its XRRP VLANs, which causes the fail-back router to stop advertising XRRP packets for the fail-over router IP addresses.)

Note

If a system operator uses **xrrp ctrl-transfer** to force a fail-back from a fail-back router to a fail-over router that has not recovered access to all of its XRRP VLANs, the fail-back is blocked and the fail-back router continues to advertise XRRP packets for the failed router's backed-up IP addresses.

- Fail-Back is triggered by an event in which all of the fail-back router's XRRP VLANs go down due to network problems or by the fail-back router rebooting. (Contrary to operation without infinite fail-back enabled, the fail-back router retains permanent control as long as at least one of its XRRP VLANs is up.)
- A system operator initiates fail-back by disabling and re-enabling XRRP on the peer router.
- Fail-Back is triggered by replacing the failed router in a protection domain with an XRRP router that is *not* configured with infinite fail-back.

Note

In a protection domain where XRRP router "A" has permanent control and peer router "B" has all of its XRRP VLANs up, replacing router "A" causes a fail-back that restores router "B" to primary control of its XRRP VLANs.

Enabling Infinite Fail-Back in a Protection Domain

As described in the chapter titled "Router Redundancy Using XRRP" in the *Advanced Traffic Management Guide* for your router, both router peers in a protection domain must have identical network access so that each can get to all the same subnets and the same end nodes without going through each other.

1. Before enabling infinite fail-back, configure XRRP on both routers in the protection domain. (Refer to the chapter mentioned in the above paragraph.)
2. Configure infinite failback on both routers in the domain.

Syntax: [no] xrrp inf-failback

Enables or disables infinite fail-back on a router running XRRP. In a given protection domain, infinite fail-back must be individually enabled on both XRRP routers to provide full fail-back control. For information on configuring general XRRP operation, refer to the chapter titled "Router Redundancy Using XRRP" in the Advanced Traffic Management Guide for your Series 5300xl router. (Default: Disabled)

3. Enable XRRP operation in the protection domain by executing the **xrrp** command in the CLI of both routers.

Initiating a Fail-Back When Infinite Fail-Back Is Enabled

Syntax: xrrp ctrl-transfer

In a protection domain where infinite fail-back has already been enabled, this command enables a system operator to manually initiate fail-back where the fail-over router has already regained access to all of its XRRP VLANs . Successful execution of this command leaves both routers in a protection domain with primary control of their configured XRRP VLANs. The command can be executed in the CLI of either router. Note that if the fail-over router has not regained access to all of its XRRP VLANs, the fail-back is blocked and the fail-back router continues to maintain permanent control of all XRRP VLANs in the domain.

Displaying the Infinite Fail-Back Configuration

Syntax: show xrrp config global

Indicates the router's global XRRP configuration, including infinite fail-back status.

```
HP ProCurve Switch 5304XL(config)# show xrrp config global

 Status and Counters - XRRP Global Configuration Information

  XRRP Enabled    : Yes
  Domain Number   : 2
  Router Number   : 1
  Failback Delay  : 10
  Infinite Failback : Enabled
```

Infinite failback status in a router configured for XRRP operation. When this shows **Enabled**, automatic fail-back does not operate and it is necessary to use the **xrrp ctrl-transfer** command to initiate a fail-back.

Note: Events in the fail-back router, such as a reboot, can also cause a fail-back. Refer to "Router Operation in the Infinite Fail-Back Mode" on page 58.

Figure 25. Example of Displaying XRRP Infinite Fail-Back Configuration Status

XRRP Log Messages

Message	Meaning
`Infinite failback has been enabled on this router.`	Infinite fail-back is enabled on the router.
`Infinite failback is not active on this router.`	Infinite fail-back has been disabled on the router.

Message	Meaning
`Peer router has permanent control`	The router is not taking primary address control because its peer has permanent (primary and secondary) address control for the protection domain.
`User has triggered failback to router` < *router-name* >	A system operator has executed the **xrrp ctrl-transfer** command to force a fail-back from the router currently having permanent (primary and secondary) control to its peer (that is, to the fail-over router).

DHCP Option 82

Introduction

The routing switch can operate as a DHCP relay agent to enable communication between a client and a DHCP server on a different subnet. Without Option 82, DHCP operation modifies client IP address request packets to the extent needed to forward the packets to a DHCP server. Option 82 enhances this operation by enabling the routing switch to append an *Option 82 field* to such client requests. This field includes two suboptions for identifying the routing switch (by MAC address or IP address) and the routing switch port the client is using to access the network. A DHCP server with Option 82 capability can read the appended field and use this data as criteria for selecting the IP addressing it will return to the client through the usual DHCP server response packet. This operation provides several advantages over DHCP without Option 82:

- An Option 82 DHCP server can use a relay agent's identity and client source port information to administer IP addressing policies based on client and relay agent location within the network, regardless of whether the relay agent is the client's primary relay agent or a secondary agent.
- A routing switch operating as a primary Option 82 relay agent for DHCP clients requesting an IP address can enhance network access protection by blocking attempts to use an invalid Option 82 field to imitate an authorized client, or by blocking attempts to use response packets with missing or invalid Option 82 suboptions to imitate valid response packets from an authorized DHCP server.
- An Option 82 relay agent can also eliminate unnecessary broadcast traffic by forwarding an Option 82 DHCP server response only to the port on which the requesting client is connected, instead of broadcasting the DHCP response to all ports on the VLAN.

Note

The routing switch's DHCP Relay Information (Option 82) feature can be used in networks where the DHCP server(s) are compliant with RFC 3046 Option 82 operation. DHCP Servers that are not compliant with Option 82 operation ignore Option 82 fields. For information on configuring an Option 82 DHCP server, refer to the documentation provided with the server application.

Some client applications can append an Option 82 field to their DHCP requests. Refer to the documentation provided for your client application.

It is not necessary for all relay agents on the path between a DHCP client and the server to support Option 82, and a relay agent without Option 82 should forward DHCP packets regardless of whether they include Option 82 fields. However, Option 82 relay agents should be positioned at the DHCP policy boundaries in a network to provide maximum support and security for the IP addressing policies configured in the server.

Option 82 Server Support

To apply DHCP Option 82, the routing switch must operate in conjunction with a server that supports Option 82. (DHCP servers that do not support Option 82 typically ignore Option 82 fields.) Also, the routing switch applies Option 82 functionality only to client request packets being *routed* to a DHCP server. DHCP relay with Option 82 does not apply to *switched* (non-routed) client requests.

For information on configuring policies on a server running DHCP Option 82, refer to the documentation provided for that application.

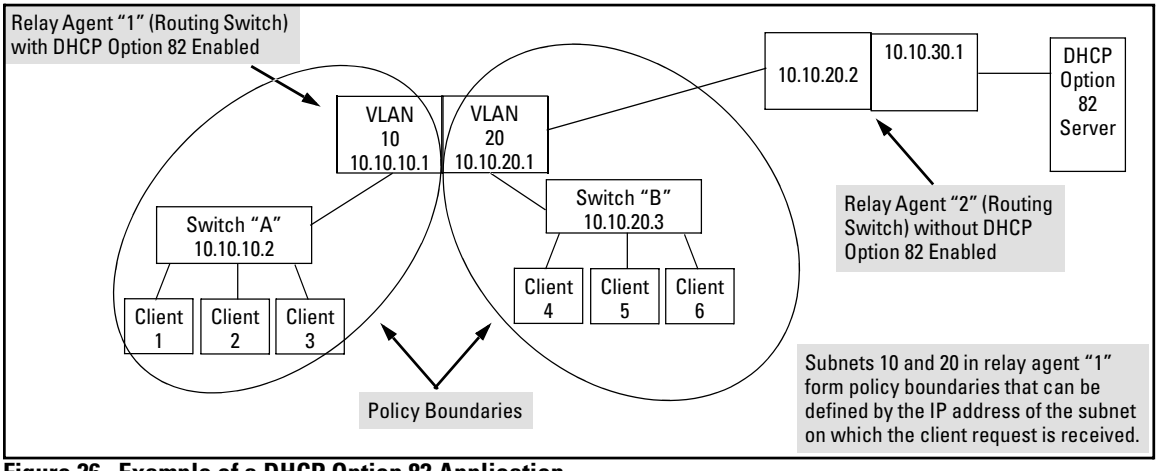

Figure 26. Example of a DHCP Option 82 Application

Terminology

Circuit ID: In Option 82 applications, the number of the port through which the routing switch receives a DHCP client request. On ProCurve fixed-port switches, the Circuit ID of a given port corresponds to the port number appearing on the front of the switch for that port. On ProCurve chassis switches, the port number for a given port corresponds to the internal ifIndex number for that port. This value is included as a suboption in an Option 82 field that the relay agent appends to a Client DHCP request before forwarding the request toward a DHCP server. (For more on Circuit ID, refer to "Circuit ID" in the bulleted list on page 65.)

DHCP Policy Boundary: For Option 82 applications, an area of a network as defined by connection to a given routing switch or subnet and/or a specific port belonging to the routing switch or subnet.

DHCP relay agent: See Relay Agent.

Forwarding Policy: The Option 82 method the routing switch uses to process incoming client DHCP requests. For a given inbound DHCP client request, the forwarding policy determines whether the routing switch will add Option 82 information, replace existing Option 82 information, or leave any existing information unchanged. The policy also determines whether the routing switch will forward the client request toward a DHCP server or drop the request. For a DHCP server response to an Option 82 client request, the routing switch can optionally perform a validation check to determine whether to forward or drop the response. Each Option 82 relay agent in the path between a DHCP client and an Option 82 DHCP server can be configured with a unique forwarding policy, which enhances DHCP policy control over discrete areas of a network.

Primary Relay Agent: In the path between a DHCP client and a DHCP server, the first routing switch (configured to support DHCP operation) that a client DHCP request encounters in the path from the client to a DHCP server.

Relay Agent: A routing switch that is configured to support DHCP operation.

Remote ID: In Option 82 applications on ProCurve switches, either the MAC address of a relay agent, or the IP address of a VLAN or subnet configured on a relay agent. This value is included as a suboption in an Option 82 field that the relay agent appends to a Client DHCP request before forwarding the request toward a DHCP server. (For more on Remote ID, refer to "Remote ID" in the bulleted list on page 65.)

Secondary Relay Agent: In the path between a DHCP client and a DHCP server, any routing switch (configured to support DHCP operation) other than the primary relay agent.

General DHCP Option 82 Requirements and Operation

Requirements. DHCP Option 82 operation is configured at the global config level and requires the following:

- IP routing enabled on the switch
- DHCP-Relay Option 82 enabled (global command level)
- routing switch access to an Option 82 DHCP server on a different subnet than the clients requesting DHCP Option 82 support
- one IP Helper address configured on each VLAN supporting DHCP clients

General DHCP-Relay Operation with Option 82. Typically, the first (primary) Option 82 relay agent to receive a client's DHCP request packet appends an Option 82 field to the packet and forwards it toward the DHCP server identified by the IP Helper address configured on the VLAN in which the client packet was received. Other, upstream relay agents used to forward the packet may append their own Option 82 fields, replace the Option 82 field(s) they find in the packet, forward the packet without adding another field, or drop the packet. (Intermediate next-hop routing switches without Option 82 capability can be used to forward—route—client request packets with Option 82 fields.) Response packets from an Option 82 server are routed back to the primary relay agent (routing switch), and include an IP addressing assignment for the requesting client and an exact copy of the Option 82 data the server received with the client request. The relay agent strips off the Option 82 data and forwards the response packet out the port indicated in the response as the Circuit ID (client access port). Under certain validation conditions described later in this section, a relay agent detecting invalid Option 82 data in a response packet may drop the packet.

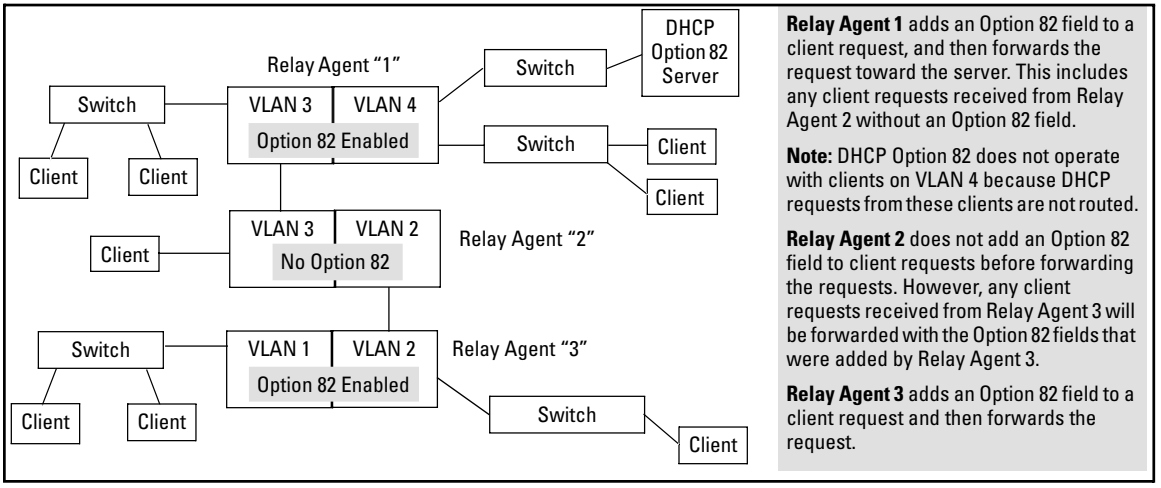

Figure 27. Example of DHCP Option 82 Operation in a Network with a Non-Compliant Relay Agent

Option 82 Field Content

The Remote ID and Circuit ID subfields comprise the Option 82 field a relay agent appends to client requests. A DHCP server configured to apply a different IP addressing policy to different areas of a network uses the values in these subfields to determine which DHCP policy to apply to a given client request.

- **Remote ID:** This configurable subfield identifies a policy area that comprises either the routing switch as a whole (by using the routing switch MAC address) or an individual VLAN configured on the routing switch (by using the IP address of the VLAN receiving the client request).
  - Use the IP address option if the server will apply different IP addressing policies to DHCP client requests from ports in different VLANs on the same routing switch.
  - Use the MAC address option if, on a given routing switch, it does not matter to the DHCP server which VLAN is the source of a client request (that is, use the MAC address option if the IP addressing policies supported by the target DHCP server do not distinguish between client requests from ports in different VLANs in the same routing switch)

    To view the MAC address for a given routing switch, execute the **show system-information** command in the CLI.

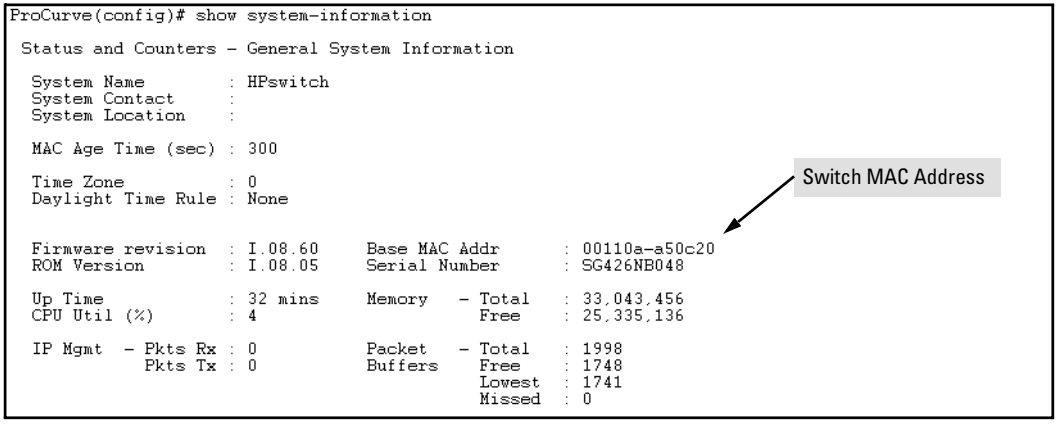

Figure 28. Using the CLI To View the Switch MAC Address

- **Circuit ID:** This nonconfigurable subfield identifies the port number of the physical port through which the routing switch received a given DHCP client request, and is necessary to identify if you want to configure an Option 82 DHCP server to use the Circuit ID to select a DHCP policy to assign to clients connected to the port. This number is the identity of the inbound port. On ProCurve fixed-port switches, the port number used for the Circuit ID is always the same as the physical port number shown on the front of the switch. On ProCurve chassis switches, where a dedicated, sequential block of internal port numbers are reserved

for each slot, regardless of whether a slot is occupied, the circuit ID for a given port is the sequential index number for that port position in the slot. (To view the Index number assignments for ports in the routing switch, use the **walkmib ifname** command.)

For example, the circuit ID for a client connected to port 11 on a ProCurve 2650-PWR (J8165A) switch is "11". However, the Circuit ID for port B11 on a ProCurve 5304xl (J4850A) is "37". (See 29, below.)

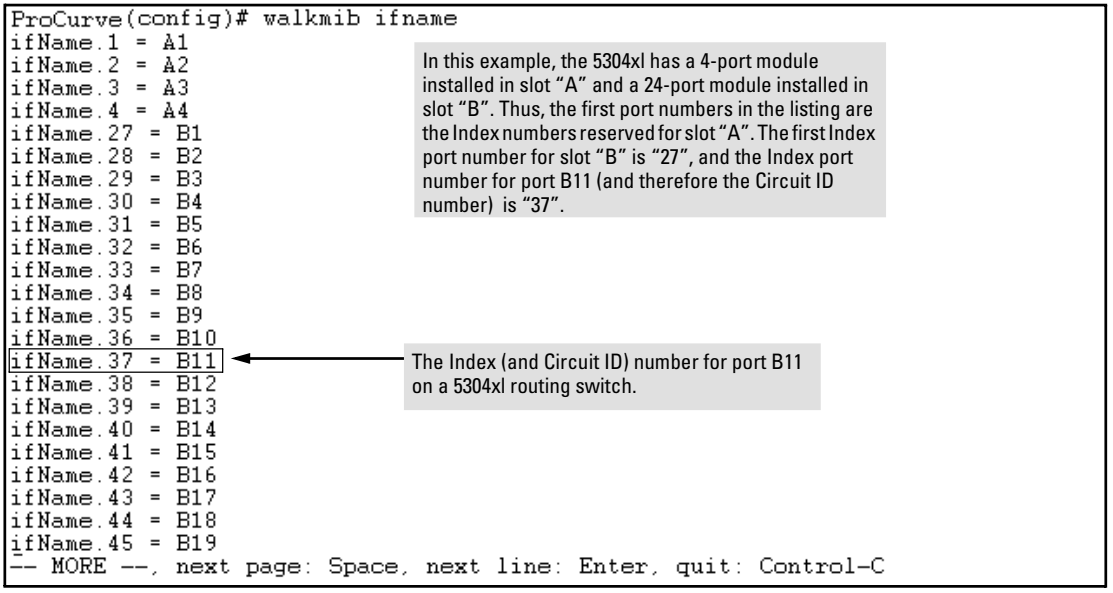

Figure 29. Using Walkmib To Determine the Circuit ID for a Port on a ProCurve Chassis

For example, suppose you wanted port 10 on a given relay agent to support no more than five DHCP clients simultaneously, you could configure the server to allow only five IP addressing assignments at any one time for the circuit ID (port) and remote ID (MAC address) corresponding to port 10 on the selected relay agent.

Similarly, if you wanted to define specific ranges of addresses for clients on different ports in the same VLAN, you could configure the server with the range of IP addresses allowed for each circuit ID (port) associated with the remote ID (IP address) for the selected VLAN.

Forwarding Policies

DHCP Option 82 on ProCurve switches offers four forwarding policies, with an optional validation of server responses for three of the policy types (**append**, **replace**, or **drop**).

Table 4. Configuration Options for Managing DHCP Client Request Packets

Option 82 Configuration	DHCP Client Request Packet Inbound to the Routing Switch	
	Packet Has No Option 82 Field	Packet Includes an Option 82 Field
Append	Append an Option 82 Field	**Append** allows the most detail in defining DHCP policy boundaries. For example, where the path from a client to the DHCP Option 82 server includes multiple relay agents with Option 82 capability, each relay agent can define a DHCP policy boundary and append its own Option 82 field to the client request packet. The server can then determine in detail the agent hops the packet took, and can be configured with a policy appropriate for any policy boundary on the path. **Note:** In networks with multiple relay agents between a client and an Option 82 server, **append** can be used only if the server supports multiple Option 82 fields in a client request. If the server supports only one Option 82 field in a request, consider using the **keep** option.
Keep	Append an Option 82 Field	If the relay agent receives a client request that already has one or more Option 82 fields, **keep** causes the relay agent to retain such fields and forward the request without adding another Option 82 field. But if the incoming client request does not already have any Option 82 fields, the relay agent appends an Option 82 field before forwarding the request. Some applications for **keep** include: • The DHCP server does not support multiple Option 82 packets in a client request and there are multiple Option 82 relay agents in the path to the server. • The unusual case where DHCP clients in the network add their own Option 82 fields to their request packets and you do not want any additional fields added by relay agents. This policy does not include the **validate** option (described in the next section) and allows forwarding of all server response packets arriving inbound on the routing switch (except those without a primary relay agent identifier.)
Replace	Append an Option 82 Field	**Replace** replaces any existing Option 82 fields from downstream relay agents (and/or the originating client) with an Option 82 field for the current relay agent.. Some applications for **replace** include: • The relay agent is located at a point in the network that is a DHCP policy boundary and you want to replace any Option 82 fields appended by downstream devices with an Option 82 field from the relay agent at the boundary. (This eliminates downstream Option 82 fields you do not want the server to use when determining which IP addressing policy to apply to a client request.) • In applications where the routing switch is the primary relay agent for clients that may append their own Option 82 field, you can use **replace** to delete these fields if you do not want them included in client requests reaching the server.
Drop	Append an Option 82 Field	**Drop** causes the routing switch to drop an inbound client request with an Option 82 field already appended. If no Option 82 fields are present, **drop** causes the routing switch to add an Option 82 field and forward the request. As a general guideline, configure **drop** on relay agents at the edge of a network, where an inbound client request with an appended Option 82 field may be unauthorized, a security risk, or for some other reason, should not be allowed.

Multiple Option 82 Relay Agents in a Client Request Path

Where the client is one router hop away from the DHCP server, only the Option 82 field from the first (and only) relay agent is used to determine the policy boundary for the server response. Where there are multiple Option 82 router hops between the client and the server, you can use different configuration options on different relay agents to achieve the results you want. This includes configuring the relay agents so that the client request arrives at the server with either one Option 82 field or multiple fields. (Using multiple Option 82 fields assumes that the server supports multiple fields and is configured to assign IP addressing policies based on the content of multiple fields.)

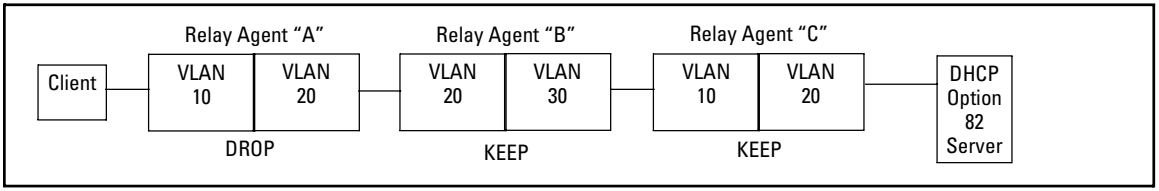

Figure 30. Example Configured To Allow Only the Primary Relay Agent To Contribute an Option 82 Field

The above combination allows for detection and dropping of client requests with spurious Option 82 fields. If none are found, then the drop policy on the first relay agent adds an Option 82 field, which is then kept unchanged over the next two relay agent hops ("B" and "C"). The server can then enforce an IP addressing policy based on the Option 82 field generated by the edge relay agent ("A"). In this example, the DHCP policy boundary is at relay agent 1.

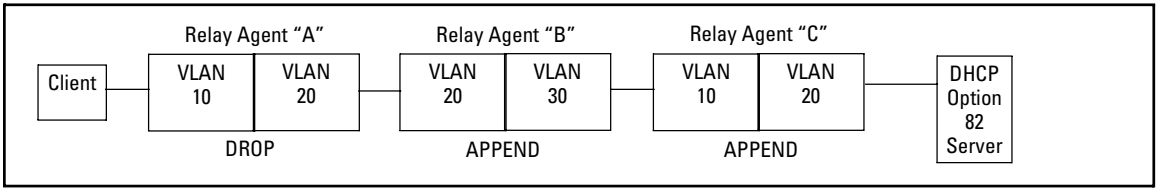

Figure 31. Example Configured To Allow Multiple Relay Agents To Contribute an Option 82 Field

This is an enhancement of the previous example. In this case, each hop for an accepted client request adds a new Option 82 field to the request. A DHCP server capable of using multiple Option 82 fields can be configured to use this approach to keep a more detailed control over leased IP addresses. In this example, the primary DHCP policy boundary is at relay agent "A", but more global policy boundaries can exist at relay agents "B" and "C".

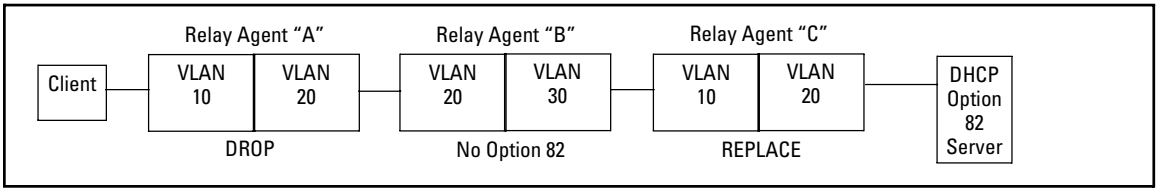

Figure 32. Example Allowing Only an Upstream Relay Agent To Contribute an Option 82 Field

Like the first example, above, this configuration drops client requests with spurious Option 82 fields from clients on the edge relay agent. However, in this case, only the Option 82 field from the last relay agent is retained for use by the DHCP server. In this case the DHCP policy boundary is at relay agent "C". In the previous two examples the boundary was with relay "A".

Validation of Server Response Packets

A valid Option 82 server response to a client request packet includes a copy of the Option 82 field(s) the server received with the request. With validation disabled, most variations of Option 82 information are allowed, and the corresponding server response packets are forwarded.

Server response validation is an option you can specify when configuring Option 82 DHCP for **append**, **replace**, or **drop** operation. (Refer to "Forwarding Policies" on page 67.) Enabling validation on the routing switch can enhance protection against DHCP server responses that are either from untrusted sources or are carrying invalid Option 82 information.

With validation enabled, the relay agent applies stricter rules to variations in the Option 82 field(s) of incoming server responses to determine whether to forward the response to a downstream device or to drop the response due to invalid (or missing) Option 82 information. Table 5, below, illustrates relay agent management of DHCP server responses with optional validation enabled and disabled.

Table 5. Relay Agent Management of DHCP Server Response Packets

Response Packet Content	Option 82 Configuration	Validation Enabled on the Relay Agent	Validation Disabled (The Default)
Valid DHCP server response packet without an Option 82 field.	**append**, **replace**, or **drop**[1]	Drop the server response packet.	Forward server response packet to a downstream device.
	keep[2]	Forward server response packet to a downstream device.	Forward server response packet to a downstream device.
The server response packet carries data indicating a given routing switch is the primary relay agent for the original client request, but the associated Option 82 field in the response contains a *Remote ID* and *Circuit ID* combination that did not originate with the given relay agent.	**append**	Drop the server response packet.	Forward server response packet to a downstream device.
	replace or **drop**[1]	Drop the server response packet.	Drop the server response packet.
	keep[2]	Forward server response packet to a downstream device.	Forward server response packet to a downstream device.
The server response packet carries data indicating a given routing switch is the primary relay agent for the original client request, but the associated Option 82 field in the response contains a *Remote ID* that did not originate with the relay agent.	**append**	Drop the server response packet.	Forward server response packet to a downstream device.
	replace or **drop**[1]	Drop the server response packet.	Drop the server response packet.
	keep[2]	Forward server response packet to a downstream device.	Forward server response packet to a downstream device.
All other server response packets[3]	**append**, **keep**[2], **replace**, or **drop**[1]	Forward server response packet to a downstream device.	Forward server response packet to a downstream device.

[1]Drop is the recommended choice because it protects against an unauthorized client inserting its own Option 82 field for an incoming request.

[2]A routing switch with DHCP Option 82 enabled with the **keep** option forwards all DHCP server response packets except those that are not valid for either Option 82 DHCP operation (compliant with RFC 3046) or DHCP operation without Option 82 support (compliant with RFC 2131).

[3]A routing switch with DHCP Option 82 enabled drops an inbound server response packet if the packet does not have any device identified as the primary relay agent (*giaddr* = null; refer to RFC 2131).

Multinetted VLANs

On a multinetted VLAN, each interface can form an Option 82 policy boundary within that VLAN if the routing switch is configured to use IP for the remote ID suboption. That is, if the routing switch is configured with IP as the remote ID option and a DHCP client request packet is received on a multinetted VLAN, the IP address used in the Option 82 field will identify the subnet on which the packet was received instead of the primary IP address for the VLAN. This enables an Option 82 DHCP server to support more narrowly defined DHCP policy boundaries instead of defining the boundaries at the VLAN or whole routing switch levels. If the MAC address option (the default) is configured instead, then the routing switch MAC address will be used regardless of which subnet was the source of the client request. (The MAC address is the same for all VLANs configured on the routing switch.)

Note that all request packets from DHCP clients in the different subnets in the VLAN must be able to reach any DHCP server identified by the IP Helper Address(es) configured on that VLAN.

Configuring Option 82 Operation on the Routing Switch

Syntax: dhcp-relay option 82 < append [validate] | replace [validate] | drop [validate] | keep > [ip | mac]

append: *Configures the routing switch to append an Option 82 field to the client DHCP packet. If the client packet has any existing Option 82 field(s) assigned by another device, then the new field is appended to the existing field(s).*

The appended Option 82 field includes the switch Circuit ID (inbound port number) associated with the client DHCP packet, and the switch Remote ID. The default switch remote ID is the MAC address of the switch on which the packet was received from the client. To use the incoming VLAN's IP address instead of the switch MAC address for the remote ID, use the* **ip** *option (below).*

replace: *Configures the routing switch to replace any existing Option 82 field(s) in an inbound client DHCP packet with one Option 82 field for the current routing switch.*

The replacement Option 82 field includes the switch circuit ID (inbound port number) associated with the client DHCP packet, and the switch remote ID. The default switch remote ID is the MAC address of the switch on which the packet was received from the client. To use the incoming VLAN's IP address instead of the switch MAC address for the remote ID, use the* **ip** *option (below).*

drop: *Configures the routing switch to unconditionally drop any client DHCP packet received with existing Option 82 field(s). This means that such packets will not be forwarded. Use this option where access to the routing switch by untrusted clients is possible.*

If the routing switch receives a client DHCP packet without an Option 82 field, it adds an Option 82 field to the client and forwards the packet. The added Option 82 field includes the switch circuit ID (inbound port number) associated with the client DHCP packet, and the switch remote ID. The default switch remote ID is the MAC address of the switch on which the packet was received from the client. To use the incoming VLAN's IP address instead of the switch MAC address for the remote ID, use the IP option (below).*

keep: *For any client DHCP packet received with existing Option 82 field(s), configures the routing switch to forward the packet as-is, without replacing or adding to the existing Option 82 field(s).*

*****For more on identifying the inbound port number, refer to "Circuit ID" in the bulleted list on page 65.*

[validate]: *This option operates when the routing switch is configured with append, replace, or drop as a forwarding policy. With validate enabled, the routing switch applies stricter rules to an incoming Option 82 server response to determine whether to forward or drop the response. For more information, refer to "Validation of Server Response Packets" on page 69.*

[ip | mac]

This option specifies the remote ID suboption the routing switch will use in Option 82 fields added or appended to DHCP client packets. The choice of type depends on how you want to define DHCP policy areas in the client requests sent to the DHCP server. (Refer to "Option 82 Field Content" on page 65.)

ip: *Specifies the IP address of the VLAN on which the client DHCP packet enters the switch.*

mac: *Specifies the routing switch's MAC address. (The MAC address used is the same MAC address that is assigned to all VLANs configured on the routing switch.) This is the default setting.*

Notes on Default Remote ID Selection: *Executing the Option 82 command without specifying either* **ip** *or* **mac** *configures the remote ID as the MAC address of the switch on which the packet was received from the client. The command options for viewing the routing switch MAC address are listed at the end of the "Remote ID" description that begins on page 65.*

Operating Notes

- This implementation of DHCP relay with Option 82 complies with the following RFCs:
  - RFC 2131
  - RFC 3046
- Moving a client to a different port allows the client to continue operating as long as the port is a member of the same VLAN as the port through which the client received its IP address. However, rebooting the client after it moves to a different port can alter the IP addressing policy the client receives if the DHCP server is configured to provide different policies to clients accessing the network through different ports.
- The IP address of the primary DHCP relay agent receiving a client request packet is automatically added to the packet, and is identified as the *giaddr* (gateway interface address). (That is, the giaddr is the IP address of the VLAN on which the request packet was received from the client.) For more information, refer to RFC 2131 and RFC 3046.
- DHCP request packets from multiple DHCP clients on the same relay agent port will be routed to the same DHCP server(s). Note that when using 802.1X on a 5300xl switch running software release E.09.xx or greater, a port's VLAN membership may be changed by a RADIUS server responding to a client authentication request. In this case the DHCP server(s) accessible from the port may change if the VLAN assigned by the RADIUS server has different DHCP helper addresses than the VLAN used by unauthenticated clients.
- Where multiple DHCP servers are assigned to a VLAN, a DHCP client request cannot be directed to a specific server. Thus, where a given VLAN is configured for multiple DHCP servers, all of these servers should be configured with the same IP addressing policy.
- Where routing switch "A" is configured to insert its MAC address as the Remote ID in the Option 82 fields appended to DHCP client requests, and upstream DHCP servers use that MAC address as a policy boundary for assigning an IP addressing policy, then replacing

switch “A” makes it necessary to reconfigure the upstream DHCP server(s) to recognize the MAC address of the replacement switch. This does not apply in the case where an upstream relay agent “B” is configured with **option 82 replace**, which removes the Option 82 field originally inserted by switch “A”.

- Relay agents without Option 82 can exist in the path between Option 82 relay agents and an Option 82 server. The agents without Option 82 will forward client requests and server responses without any effect on Option 82 fields in the packets.
- If the routing switch is not able to add an Option 82 field to a client’s DHCP request due to the message size exceeding the MTU (Maximum Transmission Unit) size, then the request is forwarded to the DHCP server without Option 82 information and an error message is logged in the switch’s Event Log.

Release E.09.21 Enhancements

Enhancement	Overview
Supports ProCurve Switch xl Access Controller Module (J8162A)	The xl Access Controller Module uses ports on a 5300xl switch to pass wired and wireless traffic to and from the network using authentication and rights administration policies from an Access Control Server 740wl or an Integrated Access Manager 760wl.

Release E.09.04 - E.09.20 Enhancements

Software fixes only; no new enhancements.

Release E.09.03 Enhancements

Software fixes only; no new enhancements.

Release E.09.02 Enhancements

The following enhancements are described in the guides referenced in each enhancement overview. For the latest version of any of these guides, refer to “Downloading Switch Documentation and Software from the Web” on page 1.

Summary of E.09.02 Enhancements

E.09.02 Enhancement	Overview
Multiple Configuration Files	This feature supports up to three different startup-config files on the switch. This enables you to easily set (and override) reboot policies, maintain different configurations for different software versions, and test new configurations while ensuring that an unattended reboot will use a proven configuration. For more information, refer to the section titled "Multiple Configuration Files on 5300xl Switches" in chapter 6 of the *Management and Configuration Guide* for your switch (January 2005 edition or later).
UDP Directed Broadcasts	This routing feature enables optional, per-VLAN UDP broadcast forwarding to support client requests sent as limited IP broadcasts to UDP application ports. For more information, refer to the section titled "UDP Broadcast Forwarding on 5300xl Switches" in chapter 11 of the *Advanced Traffic Management Guide* for your switch (January 2005 edition or later).
Virus-Throttling	This new security feature uses connection-rate filtering based on virus throttling technology, and is recommended for use on the edge of a network to detect the worm-like behavior where a host attempts to create a large number of outbound IP connections on a routed interface in a short time. The feature includes filtering and sensitivity options, and also supports the option to use specialized ACLs (Access Control Lists) to enable forwarding or dropping routed traffic from specific sources. For more information, refer to chapter 3, "Virus Throttling", in the *Access Security Guide* for your switch (January 2005 edition or later).
HTTP Support for PoE	The switch's web browser interface can now configure PoE (Power over Ethernet) operation on ports equipped for PoE operation. To access this feature, start the switch's web browser interface, click on the **Configuration** tab, and then click on the **PoE Configuration** tab.
LLDP (Link-Layer Discovery Protocol)	Enabled in the default configuration, this new feature provides a standards-based method for enabling the switches to advertise themselves to adjacent devices and to learn about adjacent LLDP devices. This feature also supports access by SNMP-based network discovery tools. Refer to the section titled "LLDP (Link-Layer Discovery Protocol" in chapter 14 of the *Management and Configuration Guide* for your switch (January 2005 edition or later).
Concurrent 802.1X and Either Web Auth or MAC Auth Per-Port	In earlier software releases, 802.1X and either Web Authentication or MAC Authentication were mutually exclusive on a port. With release E.09.02 and greater the switches allow concurrent operation of 802.1X and either Web Authentication or MAC Authentication. The combined client limit for a port configured for concurrent operation is 32. Refer to the description of the client-limit option in the section titled "1. Enable 802.1X Authentication on Selected Ports" in chapter 10 of the *Access Security Guide* for your switch (January 2005 edition or later).
Multiple 802.1X User Authentication Per-Port	This feature extends 802.1X operation to multiple-client operation and allows up to 32 authenticated clients per-port. (The client limit includes any existing clients currently authenticated through Web Authentication or MAC Authentication.) Refer to the sections titled "User Authentication Methods" and "VLAN Membership Priority" in chapter 10 of the *Access Security Guide* for your switch (January 2005 edition or later).

E.09.02 Enhancement	Overview
RADIUS IDM (Identity-Driven Management) for CoS and Rate-Limiting	The switches now allow a RADIUS server to impose 802.1p (CoS) priority and Rate Limiting on inbound traffic on a port supporting a RADIUS-authenticated client. This involves configuring vendor-specific attributes on the RADIUS server. Refer to the section titled "Configuring a RADIUS Server To Specify Per-Port CoS and Rate-Limiting Services" in chapter 6 of the *Access Security Guide* for your switch (January 2005 edition or later).
RADIUS Authorized Manager-Level Login	In the default RADIUS operation, the switch automatically admits any authenticated client to the Login (Operator) privilege level. An authenticated user authorized for the Manager privilege level must authenticate again to change privilege levels. Using the optional **login privilege-mode** command overrides this default behavior for clients having Enable (Manager) access through Telnet, SSH, the web browser interface, or the console (serial) port. (This feature does not operate with 802.1X access.) To use this feature, you must also configure the correct service-type value for the intended clients on your RADIUS server. For more information, refer to the section titled "2. Enable the (Optional) Access Privilege Option" in chapter 6 of the *Access Security Guide* for your switch (January 2005 edition or later).
RADIUS Authentication for 5300xl Web Browser Access	The **aaa authentication web < enable | login > radius** command allows the optional use of RADIUS as the primary password authentication method for the 5300xl web browser interface. Refer to the section titled "1. Configure Authentication for the Access Methods You Want RADIUS To Protect" in chapter 6 of the *Access Security Guide* for your switch (January 2005 edition or later).
802.1X Open VLAN Delay Option	To allow additional time for client authentication, the switch now provides the **aaa port-access < *port-list* > unauth-period** command. This option configures a delay of up to 255 seconds before placing a port with an unauthenticated client into the Unauthorized-Client VLAN. This gives a client with 802.1X supplicant capability more time to authenticate. Refer to the end of the section titled "1. Enable 802.1X Authentication on Selected Ports" in chapter 10 of the *Access Security Guide* for your switch (January 2005 edition or later).
Delete populated VLAN.	This feature allows you to delete a static VLAN without first having to remove all member ports from the VLAN. Any ports in the deleted VLAN that do not also belong to another static VLAN are automatically placed in the default VLAN. Refer to the section titled "CLI: Configuring Port-Based and Protocol-Based VLAN Parameters" in chapter 2 of the *Advanced Traffic Management Guide* for your switch (January 2005 edition or later).
IGMP Delayed Group Flush	This feature causes the switch to continue to filter IGMP groups for a specified additional period after IGMP leaves have been sent. This prevents unregistered traffic from being forwarded by the server during the delay period. Refer to "Configuring Delayed Group Flush", below.
Configuring Fast-Leave IGMP from the CLI	In earlier software releases, Fast-Leave IGMP was automatically enabled on all ports. Beginning with release E.09.03, you can use the CLI to disable or enable Fast-Leave IGMP on a per-port basis. (Default: Enabled.) Refer to "Configuring Fast-Leave IGMP", below.
Configuring Forced Fast-Leave IGMP from the CLI	In earlier software releases, changing the Forced Fast-Leave IGMP configuration required a **setmib** command. Beginning with release E.09.03, you can use the CLI **forcedfastleave** command to enable or disable Forced Fast-Leave IGMP. (Default: Disabled.) Refer to "Configuring Forced Fast-Leave IGMP" on page 76.

Configuring Delayed Group Flush

When enabled, this feature continues to filter IGMP groups for a specified additional period of time after IGMP leaves have been sent. The delay in flushing the group filter prevents unregistered traffic from being forwarded by the server during the delay period. In practice, this is rarely necessary on 5300xl switches, which support data-driven IGMP. (Data-Driven IGMP, which is enabled by default, prunes off any unregistered IGMP streams detected on the switch.)

Syntax: igmp delayed-flush < *time-period* >

Where leaves have been sent for IGMP groups, enables the switch to continue to flush the groups for a specified period of time. This command is applied globally to all IGMP-configured VLANs on the switch. Range: **0 - 255***; Default: Disabled (***0***).*

Syntax: show igmp delayed-flush

Displays the current **igmp delayed-flush** *setting.*

Configuring Fast-Leave IGMP

Syntax: [no] ip igmp fastleave < *port-list* >

Enables IGMP fast-leaves on the specified ports in the selected VLAN. The **no** *form of the command disables IGMP fast-leave on the specified ports in the selected VLAN. Use* **show running** *to display the ports per-VLAN on which Fast-Leave is disabled.*

Configuring Forced Fast-Leave IGMP

When enabled, Forced Fast-Leave speeds up the process of blocking unnecessary IGMP traffic per-VLAN on a switch port that is connected to multiple end nodes. (This feature does not operate on ports where the switch detects only one end node.) This command replaces the **setmib**, **getmib**, and **walkmib** commands formerly required to configure and view per-port Forced Fast-Leave.

Syntax: [no] vlan < *vid* > ip igmp forcedfastleave < port-list >

Enables IGMP Forced Fast-Leave on the specified ports in the selected VLAN, even if they are cascaded. (Default: Disabled.) The **no** *form of the command disables Forced Fast-Leave on the specified ports in the selected VLAN. Use* **show running** *to display the ports per-VLAN on which Forced Fast-Leave is enabled.*

Release E.08.53 Enhancements

Software fixes only; no new enhancements.

Release E.08.50 Enhancements

Software fixes only; no new enhancements.

Release E.08.42 Enhancements

Enhancement	Overview
Supports ProCurve Switch 10/100/1000-T xl Module (J4907A)	Provides 14 twisted-pair ports with RJ-45 connectors for 10/100/1000 Mbps (Gigabit) or 100 Mbps operation over Category 5 or better 100-ohm UTP or STP cable (category 5e recommended for Gigabit). The module also includes 2 slots for installing any of the supported ProCurve mini-GBICs.

Release E.08.30 Enhancements

Enhancement	Overview
Supports ProCurve Switch PoE xl Module (J8161A). This module is also supported in software version E.08.22 and greater.	Provides 24 10/100-TX ports that, when connected to an ProCurve 600 Redundant and External Power Supply (J8168A), can provide Power over Ethernet (PoE) power to 802.3af compliant devices.

Release E.08.07 Enhancements

Enhancement	Overview
Increased Number of IP Helper Addresses	In earlier software releases, the switch allowed up to 128 IP helper addresses to support DHCP Relay operation. Beginning with release E.08.07, you can configure up to 256 IP helper addresses. For more information on DHCP Relay and IP helper addresses, refer to the ProCurve Series 5300xl Switches *Advanced Traffic Management Guide* (part number 5990-6051, January 2005 or later). (For the latest version of this guide, refer to "Downloading Switch Documentation and Software from the Web" on page 1.)
Eavesdrop Protection when using Port Security	Using either the port-security command or the switch's web browser interface to enable port security on a given port automatically enables eavesdrop prevention on that port. This prevents the switch from using the port to flood unicast packets addressed to destination MAC addresses unknown to the switch. This blocks unauthorized users from eavesdropping on traffic intended for addresses that have aged-out of the switch's address table. (Eavesdrop prevention does not affect multicast and broadcast traffic, meaning that the switch floods these two traffic types out a given port regardless of whether port security is enabled on that port.) For more information on port security, refer to the ProCurve Series 5300xl Switches *Access Security Guide* (part number 5990-6051, January 2005 or later). (For the latest version of this guide, refer to "Downloading Switch Documentation and Software from the Web" on page 1.)

Release E.08.03 Enhancements

Software fixes only; no new enhancements.

Release E.08.01 Enhancements

Enhancement	Overview
802.1s Multiple Spanning-Tree	Adds the option for running 802.1s Multiple Spanning-Tree on the switch to enable multiple spanning-tree instances. Interoperates with legacy 802.1D (STP) and 802.1w (RSTP) spanning-tree. (Refer to the chapter titled "Spanning-Tree Operation" in the *Advanced Traffic Management Guide*—part number 5990-6051, January 2005 or later— on the ProCurve Web site.)
Protocol VLANs	Adds the capability to create layer-3 broadcast domains for IPX, IPv4, IPv6, ARP, AppleTalk, SNA, DECIat, and Netbeui protocols. (Refer to the chapter titled "Static Virtual LANs (VLANs)" in the *Advanced Traffic Management Guide*—part number 5990-6051, January 2005 or later — on the ProCurve Networking Web site.*)
PIM-DM	Adds the option to use PIM-DM to route multicast on the 5300xl switches. (Refer to the chapter titled "PIM-DM (Dense Mode)" in the *Advanced Traffic Management Guide*—part number 5990-6051, January 2005 or later — on the ProCurveNetworking Web site.)
Rate-Limiting	Provides a method for limiting the amount of bandwidth a user or device may utilize inbound on a switch port. (Refer to the chapter titled "Optimizing Traffic Flow with Port Controls, Port Trunking, and Filters" in the *Management and Configuration Guide* —part number 5990-6050, January 2005 — on the ProCurve Networking Web site.)
Guaranteed Minimum Bandwidth	Provides a method for ensuring that each of a given port's outbound traffic priority queues has a specified minimum bandwidth for sending traffic out on the link to another device. (Refer to: Chapter 10, "Optimizing Traffic Flow with Port Controls, Port Trunking, and Filters" in the Management and Configuration Guide—part number 5990-6050, February 2004 or later — on the ProCurve Networking Web site.)
Static Network Address Translation (Static NAT)	Provides the capability to conceal a "private" area of your network from the general population of users in the "public" area of the network while allowing access from the public area of the network to devices in the private area. (Refer to the chapter titled "IP Routing Features" in the *Advanced Traffic Management Guide* — part number 5990-6051, January 2005 or later — on the ProCurve Networking Web site.)
Web Authentication	Web authentication adds a new securityoption that uses a web page login to authenticate users via a RADIUS server for access to the network. (Refer to the chapter titled "Web and MAC Authentication" in the *Access Security Guide* — part number 5990-6052, January 2005 or later — on the ProCurve Networking Web site.)
MAC Authentication	MAC authentication adds a new security option that uses a device's MAC address to authenticate the device via a RADIUS server for access to the network. (Refer to the chapter titled "Web and MAC Authentication" in the *Access Security Guide* — part number 5990-6052, January 2005 — on the ProCurve Networking Web site.)

Enhancement	Overview (Continued)
MAC Lockdown/ Lockout/Secure	• MAC Lockdown enables the permanent assignment of a port or VLAN to a specific port on the switch. • MAC Lockout causes the switch to drop any traffic to or from the specified MAC address(es). • MAC Secure adds the "limited-continuous" option to the port-security command. This option sets a finite limit to the number of learned, ageable addresses (up to 32) allowed per-port. (Refer to the chapter titled "Configuring and Monitoring Port Security" in the *Access Security Guide* — part number 5990-6052, January 2005 or later — on the ProCurve Web site.)
Secure Copy and Secure FTP	Enables use of a secure, encrypted SSH session for transferring files to or from the switch. (Refer to: Appendix A, "File Transfers" in the *Management and Configuration Guide*—part number 5990-6050, January 2005 or later — on the ProCurve Networking Web site.)
Front-Panel Security	Provides the option for enabling or disabling some of the functions of the Reset and Clear buttons on the switch's front panel. This feature also provides the ability to disable password recovery for situations requiring a higher level of security. (Refer to the chapter titled "Configuring Username and Password Security" in the *Access Security Guide* — part number 5990-6052, January 2005 or later — on the ProCurve Networking Web site.)
RIP Debug Logging Logging Facility	• RIP Debug Logging adds RIP event logging to the switch's debug destination options. • Logging Facility adds the capability for specifying the destination subsystem the configured SyslogD servers must use. (The default is the "user" subsystem.) (Refer to: Appendix C, "Troubleshooting" in the *Management and Configuration Guide* — part number 5990-6050, January 2005 or later — on the ProCurve Networking Web site.)
Eight-Port Trunking	Increases the number of ports allowed in a trunk from four to eight. (Refer to the chapter titled "Optimizing Traffic Flow with Port Controls, Port Trunking, and Filters" in the *Management and Configuration Guide* — part number 5990-6050, January 2005 or later — on the ProCurve Networking Web site.)
Auto MDI-X	Provides CLI commands for changing the cable-configuration support on the switch's copper ports. The options include auto-MDIX (the default), MDI, and MDI-X. (Refer to the chapter titled "Optimizing Traffic Flow with Port Controls, Port Trunking, and Filters" in the *Management and Configuration Guide* — part number 5990-6050, January 2005 or later — on the ProCurve Networking Web site.)
Port Duplex Mismatch Detection	Provides automatically enabled duplex mismatch detection on all ports and generates an Event Log message for detected mismatches.
Flow sampling with sFlow	Adds sFlow as a flow-sampling method for use with applicable network management software. (Refer to the documentation provided with your network management software.)

Release E.07.40 Enhancements

Includes the same Boot ROM update as Release E.07.37.

Release E.07.37 Enhancements

Boot ROM Update, Version E.05.04 — Modifications have been made to Boot ROM to allow for a larger flash memory space to be available for future software releases and switch feature enhancements. The updated Boot ROM is backwards compatible with older software. Downloading E.07.37 will automatically update Boot ROM code upon switch reboot.

You can confirm that the Boot ROM has been updated using the **show flash** command. See the example below:

```
ProCurve# show flash
Image            Size(Bytes)   Date   Version
-----            -----------  -------- -------
Primary Image    : 2901599    09/26/03 E.07.37
Secondary Image  : 2784788    08/13/03 E.07.34
Boot Rom Version: E.05.04
Current Boot     : Primary
```

Release E.07.34 Enhancements

Software fixes only; no new enhancements.

Release E.07.30 Enhancements

Software fixes only; no new enhancements.

Release E.07.29 Enhancements

Software fixes only; no new enhancements.

Release E.07.27 Enhancements

Software fixes only; no new enhancements.

Release E.07.22 Enhancements

Software fixes only; no new enhancements.

Release E.07.21 Enhancements

To Locate Publications Supporting E.07.21 Features:

1. Go to HP's ProCurve web site at **http://www.procurve.com**.
2. Click on **technical support**, then **manuals**.
3. Click on the name of the product for which you want documentation.
4. Select the document indicated in the enhancement description (6) for the desired feature.

(HP recommends periodically visiting the ProCurve Networking Web site to keep up-to-date with the latest documentation available for ProCurve Series 5300xl switch products.)

Table 6. Release E.07.21 Enhancements

Enhancement	Overview
Access Control Lists (ACLs)	Layer 3 IP filtering with ACLs lets you improve network performance and restrict network use by creating policies for switch management access and application access security.[1]
Debug and Syslog Messaging Operation	These features provide a method for recording messages you can use to help in debugging network-level problems such as routing misconfigurations and other protocol details.[1]
SNMPv3	The Series 5300xl switches now support SNMPv3 to enhance the security of SNMPv3 traffic. It include authentication and/or encryption of Management traffic configurable at the operators discretion.[1]
Meshing improvements	The Series 5300xl switches now have improved meshing features. They include greater configuration checks for meshes with (only) Series 5300xl and backwards compatibility mode for reduced connect times with legacy meshing devices.[1] **Note:** If you update to release E.07.21 or later from a code version earlier than E.07.21, *and* if the Series 5300xl switch belongs to a switch mesh domain that includes any ProCurve 1600M, 2400M, 2424M, 4000M, or 8000M switches, then you must execute the backward compatibility mode command (mesh backward-compat). Otherwise, the Series 5300xl switch will not connect to the mesh.
OSPF Authentication	Adds MD5 encryption for authenticating OSPF packets. Encryption keys are managed by a centralized Key Management System (KMS).[1]
SSHv2	Updates SSH to support SSHv2. This allows for the use of PEM encoded keys and greater compatibly to SSH client software.[2]
SSL	The Series 5300xl Switches now support Secure Socket Layer transactions for Web management access. This allows the switch to authenticate itself to the user and to establish a secure connection. There is support for self-signed and CA signed certificates to allow the administrator to choose the level of security required.[2]
XRRP	The feature used by the ProCurve Series 5300xl switches to provide router redundancy or failover – a backup router in case one fails. XRRP is similar to the industry standard VRRP (Virtual Router Redundancy Protocol), although the details of the operation are different.[1]
IGMPv3	Adds support for the IGMPv3 Join request.[1]

[1] Refer to the *Management and Configuration Guide for the ProCurve Series 5300xl Switches*, part number 5990-6050, January 2005 or later, on the ProCurve Networking Web site.

[2] Refer to the *Access Security Guide for the ProCurve Series 5300xl Switches*, part number 5990-6052, January 2005 or later, on the ProCurve Networking Web site.

Release E.06.10 Enhancements

Adds support for the J4852A ProCurve Switch xl 100-FX MTRJ module. Refer to the *ProCurve Switch xl Modules Installation Guide* (part number 5990-6076, May 2004 or later).

Release E.06.05, E.06.03, and E.06.02 Enhancements

Software fixes only; no new enhancements.

Release E.06.01 Enhancements

To Locate Publications Supporting E.06.01 (and greater) Features:

1. Go to HP's ProCurve web site at **http://www.procurve.com**.
2. Click on **technical support**, then **manuals**.
3. Click on the name of the product for which you want documentation.
4. Select the publication indicated in the enhancement description (below) for the desired feature.

(HP recommends periodically visiting the ProCurve Networking Web site to keep up-to-date with the latest documentation available for ProCurve Series 5300xl switch products.)

Enhancement	Overview
HP J4860A LH-LC Mini-GBIC	New long-haul mini-GBIC support for Series 5300xl switches.[1]
New Flow Control Command	The Series 5300xl switches enable per-port flow control. Beginning with release E.06.0x, use the (global) Flow Control command to enable flow control on the switch, then enable flow control on the desired port(s).[2] **Note:** If you have enabled flow-control on individual ports while using software version E.05.04, but then downloaded software version E.06.01 (or greater) and rebooted the switch, flow control will be disabled globally on the switch (the default) and therefore will not operate on the individual ports previously configured to allow flow control. To resume the configured per-port flow-control activity, you must enable global flow control.
Change in Default for RIP Redistribution of Connected Routes	Formerly, RIP redistributed both static and connected routes by default. Now, the factory-default RIP operation automatically redistributes connected routes, but not static routes.[2]
Change in Default State DHCP-Relay and Helper-Addresses	Changes the factory-default state to now enable DHCP-Relay. You can determine the current state and list Helper addresses.[2] **Note:** With DHCP-Relay disabled, if you update from release E.05.04 to E.06.01 or greater, then reboot the switch, DHCP-Relay becomes enabled. (However, to use DHCP-Relay, you will still need to configure IP Helper addresses.)
Trace Route	Provides a new feature for tracking the path of a packet between the switch and a destination IP address.[2]

Enhancement	Overview (Continued)
IP Address Command Change	This change simplifies multinetting on VLANs.[2]
802.1x Open VLAN Mode (Authorized-Client and Unauthorized-Client VLANs)	Provides more flexibility for authenticating clients lacking 802.1x supplicant software, and an additional provision for controlling VLAN access by authenticated clients.[3]

[1] Refer to the *Switch xl Modules Installation Guide, part number 5990-6076, May 2004 or later.*
[2] Refer to the *Management and Configuration Guide for the ProCurve Series 5300xl Switches*, part number 5990-6050, January 2005 or later.
[3] Refer to the *Access Security Guide for the ProCurve Series 5300xl Switches*, part number 5990-6052, January 2005 or later.

Software Fixes in Release E.06.*xx* through E.10.*xx*

Release E.05.04 was the first software release for the ProCurve Series 5300xl switches.

Release E.10.23

Problems Resolved in Release E.10.23

- **CLI/DHCP (PR_1000286898)** — Under some conditions, the CLI may freeze or lock up when the DHCP relay agent is configured.
- **Crash (PR_1000307280)** — Inconsistent or incorrect STP data may cause the switch to crash with a message similar to:

  ```
  Software exception at stp_mib.c:248 -- in 'mSnmpCtrl', task ID =
  0x12d14b8\n-> ASSERT:  failed.
  ```

- **IGMP (PR_1000301557)** — Data-driven IGMP does not prevent flooding when no IP address exists on a VLAN.
- **SNMP (PR_1000295753)** — Removing 'public' SNMP community generates an empty Event Log message.
- **IP Forwarding (PR_1000305739)** — When a user attempts to configure 'ip forward-protocol netbios-dgm', the switch incorrectly configures 'ip forward-protocol netbios-ns' instead.
- **RSTP (PR_1000306227)** — RSTP TCNs cause high CPU utilization and slow software based routing.

Release E.10.22 (Never released)

Problems Resolved in Release E.10.22

- **Event Log (PR_1000306769)** — When an OS upgrade causes an FEC trunk to be converted, the following messages are logged:

  ```
  [datestamp] mgr: Config file converted due to OS upgrade
  W [datestamp] mgr: Unsupported feature "FEC" for trunk configuration;
  see release notes
  ```

- **Event Log/ARP (PR_1000293466)** — Generic Link Up message not showing up and unnecessary flushing of ARP cache.
- **LLDP (PR_1000301069)** — When LLDP admin status of a port changes from TX to DIS/RX, the switch does not always send out shutdown frames.

- **LLDP (PR_1000303500)** — Missing LLDP-MED information when using command: `"show lldp info remote-devices"`.
- **Meshing (PR_1000300756)** — Time delay in switch when reporting a mesh link being down.
- **Web Authentication (PR_1000302945)** — When a client fails authentication and is assigned to the Unauthorized VLAN, it cannot communicate with other clients on the Unauthorized VLAN.

Release E.10.21 - Never released

Release E.10.20

Problems Resolved in Release E.10.20

- **Key Management System (PR_1000287934)** — Some Key Management System (KMS) configuration commands have no effect.

Releases E.10.11 to E.10.19 were never built.

Release E.10.10

Problems Resolved in Release E.10.10

FEC (PR_1000281715) — Switch has no FEC support but shows FEC information in help text.

Release E.10.09

Problems Resolved in Release E.10.09

- **Config (PR_1000301498)** — The user cannot manually configure an IP address using the "setup" menu.
- **FEC/CDP (PR_1000285111) —** FEC and CDP transmit removal.
- **Routing (PR_1000297773) —** Certain types of traffic cause the switch to route very slowly and drop packets.
- **RSTP (PR_1000297195)** — The switch repeatedly flushes its MAC address table, resulting in intermittent flooding of all traffic.

Release E.10.08

Problems Resolved in Release E.10.08

- **Enhancement (PR_1000290489)** — Support for “Friendly Port Names” was added.

Release E.10.07

Problems Resolved in Release E.10.07

- **802.1X (PR_1000290453)** — 802.1X stops and restarts the accounting session during re-authentication.
- **802.1X (PR_1000216987)** — An 802.1X client may age out prematurely if it communicates in multiple VLANs.
- **802.1X (PR_1000235378)** — When client based authentication was introduced in E.09.02, the port based authentication mode, which allows an unlimited number of clients per port, was inadvertantly removed
- **Crash (PR_1000290428)** — When a non-genuine mini-GBIC is installed into the switch, the switch may crash with a message similar to:

  ```
  "chassis: Slot A Software exception at port_sm.c:316 -- in
  'mPmSlvCtrl', task ID = 0x4059c9d4."
  ```

- **Web-Authentication (PR_1000230444)** — Some clients may not receive a Web-Authentication screen when using port-based Web-Authentication. This may occur if a client receives the same unauthorized DHCP address that a previous authorized client had used.

Release E.10.06

Problems Resolved in Release E.10.06

- **RSTP (PR_1000286883)** — Slow RSTP fail-over and fall-back time.

Release E.10.05

Problems Resolved in Release E.10.05

- **ACL (PR_1000283338)** - The commands "show port-access mac" and "show port-access web" incorrectly display the number of clients authenticated.
- **Crash (PR_1000282444)** - When enabling OSPF MD5, the switch may crash with a message similar to:

```
Software exception at exception.c:373 -- in 'mSess1'.
```

- **mini-GBIC (PR_1000283081)** — After hot-swapping a mini-GBIC, the Link and Activity LEDs do not turn on.
- **mini-GBIC (PR_1000283082)** — Some Gigabit LX mini-GBICs may fail when the mini-GBIC switch module is hot-swapped.
- **mini-GBIC (PR_1000283084)** — When a mini-GBIC is removed from the module, the Fault and Port LEDs will continue to flash.
- **RADIUS (PR_1000285456)** — If more than one RADIUS assigned vendor specific attribute (including Port-cos, rate-limiting-ingress, or ACLs) is configured with a non-vendor specific attribute, only the first vendor specific attribute may be recognized by switch.
- **Web UI/mini-GBIC (PR_1000279145)** — When using the web user interface, the switch will not display an indication of the Gigabit 1000Base-T mini-GBIC (J8177B) from the Configuration tab "Device View".
- **XRRP (PR_1000280213)** — When configuring a XRRP instance, the following error message is logged, although the particular subnet is configured properly

```
No subnet configured for the IP address.
```

Release E.10.04

Problems Resolved in Release E.10.04

- **Console/TELNET (PR_1000278912)** - The 5300xl console will lock up when connected via the console port and attempting to establish a TELNET connection into a remote switch.
- **Meshing (PR_1000218463)** - If a mesh link goes down and a redundant (xSTP) link external to the mesh goes into a forwarding state, connectivity across the mesh may be lost for a previously learned MAC address.
- **SNMP (PR_1000003378)** - SNMP switch time may drift with event log updates occurring every 1.5 hours.

Release E.10.03

Problems Resolved in Release E.10.03

- **MAC Auth/Web Auth (PR_1000244293)** - Web and MAC authentication clients do not de-authenticate immediately.
- **Config (PR_1000246102)** - The **show config** command indicates a configuration file named "config" already exists.

Release E.10.02

Problems Resolved in Release E.10.02

- **CLI (PR_1000223516) —** CLI hang when performing command involving 802.1X, Web/MAC authentication or port.
- **Config (PR_1000207697) —** Loading a startup-config file fails when file declares a new VLAN as a management VLAN.
- **Config (PR_1000215370) —** Configuration file upper/lower case is not consistent. When looking and viewing file there is inconsistencies between what is shown and what can be tab completed
- **Crash (PR_1000229613) —** A secondary flash update via PCM+ causes a bus error crash.
- **Crash (PR_1000243402) —** Null semaphore usage in SSH. (The switch may crash when "exit" is issued from slot context.)
- **Crash (PR_1000233993) —** A switch crash occurs after an **snmpgetnext** on the CDP MIB. Software exception at `exception.c:373 -- memory system error`.
- **Crash (PR_1000232283) —** Multiple TFTP requests from PCM cause a switch crash "`Software exception at fileTransferTFTP.c:182 -- in 'mftTask', task ID = 0x107ee0`.
- **Crash-OSPF (PR_1000234773) —** Within a VLAN configured with an OSPF key-chain 255, any time an external device is plugged into the VLAN on the 5300xl switch configured with the key-chain, the 5300xl switch crashes with an `ifInfo task: SubSystem 0` indicator.
- **J8162A (PR_1000219468) —** No Event Log message when user reboots J8162A Access Control Module without first shutting it down.
- **LLDP (PR_1000220937) —** LLDP advertises the base MAC address when no VLAN-IP exists. LLDP advertises "127.0.0.1" as the management TLV information on a port when no IP address is configured on any VLAN that this port belongs to. It should advertise the switch's base MAC address instead.
- **LLDP (PR_1000241315) — show lldp** issues:
  - Port descriptor may be corrupted, as displayed, if > 4.
  - PortID type of MAC-address is truncated.
  - ChassisId of type network-address is shown in MAC address format.
  - Remote Management Address type ethernet is shown in IP format.
  - Inconsistent name for PortDescr in detail view and summary ("PortDesc" vs "Port-Name").

- **MST (PR_1000227432) —** Learning flag is not set when CIST port states are transitioning.
- **Other (PR_1000214324) —** J8162A Access Controller module VLAN base configuration record. Should not create an "access-controller vlan-base" command in the remote configuration file if there is no J8162A blade in the system AND no J8162A has been configured for the switch. (There are no client VLANs on the switch.)
- **Other (PR_1000085508) —** A mini-GBIC is not recognized if the J4878A is hot-swapped during boot-up.
- **Other (PR_1000221089) —** The 64-bit counters are not correct.
- **Other (PR_1000227607) —** Problem with **show fault-finder**. The table contains two extra empty IDs.
- **Other (PR_1000235094) —** With HTTP/RADIUS, a username/password box appears for every switch between the manager and operator pages. If the Web user interface for the switch management is configured login/enable with either RADIUS/local or local/RADIUS, and local username/passwords are set and are not the same as for RADIUS, then a username/password box/prompt appears for every instance where there is a switch between an operator-level Web page (such as Status) and a manager-level Web page (such as Configuration), and the reverse.
- **PIM-DM (PR_1000235581) —** PIM DM does not always prune when Switch receives a PIM Prune message.
- **Port Security (PR_1000244293) —** Web/MAC Authentication clients do not de-authenticate immediately.
- **RMON (PR_1000240752) —** The RMON and FFI severities need correct mapping. The FFI severity levels are from low to high, whereas the RMON severity levels are mapped from high to low.
- **SFTP (PR_1000227950) —** SFTP image "puts" to a switch low on memory does not succeed. The Event log shows

  ```
  update: Disabled RMON to retrieve memory for download
  ```

  on a 5300xl switch that has ~6.7M of free memory available. The transfer does not take place and the Event log message is displayed for every attempt.
- **STP (PR_1000234771) —** The switch does not do spanning-tree fast-aging when Web-authentication changes aging for LPORT.
- **Update (PR_1000227992) —** SFTP allows an image upload of firmware for a different platform (switch model).
- **Virus Throttling (PR_1000237928) —** Add port names to the rest of the virus throttling RMON messages. Three of the existing virus throttling messages do not have the LPORT information.

- **XRRP/802.1s (PR_1000240958) —** XRRP fail-over communication issues when MSTP is also configured.

Release E.09.29 (Beta Only)

Problems Resolved in Release E.09.29

- **Crash (PR_1000229656) —** When RADIUS server is unavailable, the following message appears:

```
Software Exception at exception.c:373 -- in 'tHttpd', task ID =
0x257dda8 -> Memory system error at 0x 24ea750 - memPartFree
```

- **Crash (PR_1000235856) — show tech**' causes:

```
Software exception at dmaRx.c:868 -> ASSERT
```

- **Other (PR_1000221018) —** Menu leaves proxy-ARP configured when IP routing is disabled.
- **Other (PR_94943) —** The Setup screen allows an illegal configuration (Proxy-ARP). Using the "Setup" utility, you can toggle the Proxy-ARP entry (at the bottom of the screen) even though IP routing is NOT enabled on the system.
- **Proxy ARP (PR_94943) —** Setup screen allows illegal configuration (proxy-arp).
- **Proxy ARP (PR1000221018) —** Menu leaves proxy-arp configured when routing is disabled.
- **XRRP (PR_1000217922) —** XRRP router in infinite-failback mode can sometimes give up control of its IP address.

Release E.09.26 (Beta Only)

Problems Resolved in Release E.09.26

- **Config (PR_1000228888) —** The console becomes unresponsive ("hangs" or "freezes") when attempting to issue a configuration command, and then 802.1X and Web/MAC Authentication functions in the Switch do not operate.
- **Config (PR_1000229407) —** Edge ports on a switch with MSTP are lost when the configuration is TFTPed in from a TFTP server.
- **Hang (PR_1000228888) —** The Console becomes unresponsive ("hangs" or "freezes") when attempting to issue a configuration command, resulting in 802.1X and Web/MAC Authentication functions in the Switch ceasing to operate.

- **MSTP (PR_1000229407) —** The Switch loses the MSTP 'edge-port' configuration when the user TFTPs the configuration file from a server.

Release E.09.25 (Beta Only)

Problems Resolved in Release E.09.25

- **Config (PR_1000233062) —** Download of Configuration to alternate configuration not working.
- **XRRP (PR_1000217922) —** There is a small possibility that the XRRP Router will fail back to the XRRP peer even if infinite failback is enabled when running 802.1d and XRRP routers are redundantly connected to a large switch domain.

Release E.09.24 (Beta Only)

Problems Resolved in Release E.09.24

- **XRRP (PR_1000217922) —** XRRP router in infinite-failback mode can sometimes give up IP address control.

Release E.09.23 (Beta Only)

Problems Resolved in Release E.09.23

- **802.1s (PR_1000207608) —** After the Spanning Tree Root Bridge is negotiated the non-root ProCurve Switch continues to send out BPDUs claiming to be the Spanning Tree Root, resulting in possible instability in the STP topology. Support: This is the 'Force10/yahoo' fix, merged from the 2800.
- **Config (PR_1000215024) —** Memory leak when loading a configuration file from a TFTP server.
- **MST (PR_1000222230) —** MSTP (802.1s) sometimes fails to block a loopback connection.
- **Web UI (PR_1000214188) —** Problems with scroll bar after resizing window.
- **Web UI (PR_1000223183) —** VLANs are not displayed in QoS configuration.
- **Web (PR_1000214188) —** Problems with the scroll bar after resizing window.
- **Web (PR_1000223183) —** VLANs are not displayed in QoS configuration screen.

Release E.09.22

Problems Resolved in Release E.09.22

- **CLI (PR_1000223516)** — CLI hangs when entering certain port commands such as those involving Web MAC authentication or 802.1X.
- **MDI/MDI-X (PR_1000220687)** — Switch does not report the state of MDI/MDI-X correctly for ports on the J8161A PoE module.
- **RSTP (PR_99049)** — Switch does not detect and block network topology loops on a single port. For example, the port connects to a hub that has a loop or the port connects to an inactive node via IBM 'Type 1' cable.

Release E.09.21 (Beta Only)

Problems Resolved in Release E.09.21

- **CLI/GVRP (PR_1000216305)** — The GVRP command **no VLAN <*vid*> forbid <*ports*>** incorrectly deletes ports configured for AUTO mode.
- **Crash (PR_1000216170)** —The switch will crash with an 'mftTask' Bus Error after uploading a startup-configuration from a TFTP server. The switch accepts the command with no errors. However, the system will immediately crash after the reboot.
- **Crash (PR_1000021764)** —The switch may crash with a message similar to:

  ```
  Software exception in dmaRx.c:839.
  ```

- **Crash/LLDP (PR_1000217480)** —The switch may crash with a Bus error specifying "Task = mlldpCtrl".
- **Crash/SSH (PR_1000192010)** —The switch may crash with a message similar to:

  ```
  Software exception at exception.c:328 -- in 'tSsh0', task ID = 0x101c590.
  ```

- **Crash/Static Route (PR_1000217354)** —The switch may crash with a `Bus error in mSnmpCtrl` when adding a less-specific static route.
- **LLDP (PR_1000202129)** — The command **show lldp info remote** does not provide correct information.
- **LLDP/Mesh (PR_1000216041)** — Switch does not issue an Event Log message if LLDP is configured inconsistently among mesh neighbors.
- **MAC Authorization (PR_1000212868)** — MAC Authorization ages out a client prematurely when the client passes traffic in multiple VLANs.

- **Port Security (PR_1000210932)** — Open VLAN mode (Unauthorized VLAN) does not work correctly with any port-security learn-mode.
- **SSH (PR_1000207275)** — The Codenomicon test tool causes memory leaks in SSH.
- **Virus Throttling (PR_1000213532)** — The command **show conn throttled-hosts** displays hosts on ports set to notify-only.
- **Web UI (PR_1000191635)** — Port column may not be sorted correctly in all Web user interface screens.

Release E.09.10 (Not a General Release)

Problems Resolved in Release E.09.10

- **OSPF/Routing (PR_1000202847)** — Asymmetrical routing with equal-cost paths results in high CPU utilization and dropped packets. NOTE: This bug fix is NOT included in E.09.21, but it is in releases E.09.22 and later.

Release E.09.09 (Beta Only)

Problems Resolved in Release E.09.09

- **XRRP (PR_1000217651)** — XRRP may cause excessive event log messages.

Release E.09.08 (Beta Only)

Problems Resolved in Release E.09.08

- **Crash (PR_1000207542)** — The switch may crash with a bus error or a task hang.
- **Port Security (PR_1000203984)** — Switch allows a user to enter more MAC addresses than the configured limit.

Release E.09.07 (Beta Only)

Problems Resolved in Release E.09.07

- **QoS (PR_1000216179)** — QoS DSCP is not maintained when the switch routes the packet.

Release E.09.06 (Beta Only)

Problems Resolved in Release E.09.06

- **Config/Stack (PR_1000216051)** — Copying a previously saved startup-configuration that has "stack join (mac address)" to a member switch of the IP stack breaks the membership of that same stack. Stack commander reports member "mismatched".
- **Web (PR_80857)** — Java files are JDK 1.1, which are not Win2k compliant. (For this fix they were recompiled using JDK 1.2.)
- **Web UI/Port Status (PR_93721)** — The Port Status screen does not display all ports in the Web user interface, and the scroll bar does not work.

Release E.09.05 (Beta Only)

Problems Resolved in Release E.09.05

- **CLI/STP (PR_1000214598)** — Switch does not accept the "spanning-tree 1 mode fast" CLI command. Switch does accept and implement the span tree port mode fast setting via the menu options. However, the setting does not show up in the running configuration.
- **LLDP (PR_1000213942)** — Neighbor entry is deleted and re-learned when port admin-status is changed from rxonly to tx_rx.

Release E.09.04 (Beta Only)

Problems Resolved in Release E.09.04

- **PIM (PR_1000206791)** — With PIM enabled, an IGMP "leave" received from one subscriber causes all IP multicast streams to pause and then resume.

Release E.09.03

Problems Resolved in Release E.09.03

- **SNMP Trap (PR_1000212170) —** Switch transmits Warm and Cold Start traps with an agent address of 0.0.0.0.
- **Telnet Hang (PR_1000215388) —** When a user executes the **show configuration <*filename*>** command in a Telnet session and the file is longer than a single screen, the user's Telnet session may hang.

Release E.09.02 (Beta Only)

Problems Resolved in Release E.09.02

- **DMA Driver (PR_1000209595) —** ASSERT_RESOURCE prints `No resources available` if it fails.
- **IP Addmgr (PR_1000202539) —** ARP cache gets cleared whenever a port comes up.
- **IP Addmgr (PR_1000206356) —** Software exception at `ipamMAcl.c:712` -- host table filled with no ACLs.
- **MGR (PR_1000202237)** — VLAN MAC table flushing does not work.
- **Mirror Port (PR_1000204834) —** Mirror-Port adds a VLAN tag to untagged, monitored traffic.
- **NAT (PR_1000199309) —** NAT getting lost when cable moved.
- **Other (PR_1000204617) —** Port mirroring and ACLS cause blade assert at dmaRx.c:1319.
- **Other (PR_1000208358) —** Mac-to-Host route table mismatch.
- **Other (PR_1000092011) —** Software exception at `c:356 -- in 'mHttpCtrl'.`
- **Password (PR_1000201614) —** Non-Null terminated password causes bus error crash in setup menu.
- **Rate Limiting (PR_1000201978) —** Radius rate-limiting-ingress should allow greater than 100%.
- **RMON (PR_1000011690)** — When RMON thresholds in the switch are exceeded, no trap is generated.
- **Self-Test (PR_1000200371) —** Ports are not isolated during the selftest internal loopback testing.
- **SNMP (PR_1000196170) —** Traps are not buffered before the IP stack is initialized, causing the possibility of missing some traps generated during startup.
- **SNTP (PR_1000199632) —** NTP (SNTP) version 4 broadcast ignored by switch.
- **Tst.System (PR_1000204782) —** Bus error when copying a configuration to the switch.

Release E.08.53

Problems Resolved in Release E.08.53

- **IP Helper/DHCP Relay (PR 1*197046)** — May not handle "DHCP Inform" relay properly.
- **NAT (PR_1*199309)** — Routing to some end nodes fails when a cable is moved from one port to another or when the equivalent action happens due to XRRP fail-over or fail-back.

- **NAT (PR_1*203787)** — NAT problem when the switch has multiple VLANs configured on a port with routing enabled (that is, the one-armed router scenario)
- **SNMP/Authorized Manager (PR_1*86062)** — SNMP Sets allowed when in Operator mode and IP Authorized-Manager is set.

Release E.08.50

Problems Resolved in Release E.08.50

- **Agent (PR_1000091445)** — The switch will not respond to incoming packets, which are addressed to the switch, that have a source Ethernet address of all 0s.
- **Auto-TFTP (PR_1000187649)** — After being disabled, Auto-TFTP will not allow a forced download of an OS.
- **Crash (PR_Unknown)** — The switch may crash with a message similar to:
  `Software exception at dmaTx.c:313 -- in 'tSvcWorkQ'`
- **DHCP Relay (PR_1000188635)** — The DHCP Relay agent sometimes forwards DHCP packets without modifying the MAC SA (that is, without inserting the 5300's MAC address into the SA field of the packet).
- **DHCP Relay/XRRP (PR_1000189359)** — When DHCP Relay is relaying a packet from the client to the server, it is using its router MAC address instead of XRRP MAC address as source MAC if the IP Helper address is a broadcast IP address.
- **Meshing (PR_1000193662)** — Packets with an all-zeroes MAC Destination address are not forwarded through a mesh.
- **Meshing / Address Manager (PR_1000192016)** — Broadcasts sent into a mesh are not always delivered.
- **NAT/XRRP (PR_1000188709)** — The switch sometimes does not use the XRRP virtual MAC address.
- **Performance (PR_1000188694)** — All IP Multicast packets are routed by the CPU resulting in slower performance for any CPU-dependent activity.
- **Web Authentication (PR_1000189519)** — The Web Authentication URL has an extra slash: "http:///webauth/index.html". As a result, the SSL authentication page does not come up correctly

Release E.08.42

Problems Resolved in Release E.08.42

- **ACL (PR_1000023119)** — An invalid VLAN ACL will remain in the configuration.
- **CLI (PR_1000002138)** — Incorrect message displayed in the CLI **aaa port-access** command.
- **CLI (PR_1000022443)** — Within the CLI Menu context, user unable to set a port as an untagged member of a VLAN.
- **CLI (PR_1000085477)** — The word "Specify" in 'ip route' is misspelled.
- **CLI (PR_1000085495)** — The word "unavailability" is spelled wrong for the "radius server dead-time" description within the CLI.
- **Crash (PR_1000021489)** — The switch may crash with a message similar to:
  `Software exception at i2cdriver.c:75 in 'swInitTask'`
- **Crash (PR_1000021567)** — The switch may crash with a message similar to:
  `Software exception @ ipaddrmgrSCtrl.c:565`
- **Crash (PR_1000022106)** — The switch may crash with a message similar to:
  `Exception hit in alphaSLaveLearn.c:1534`
- **Crash (PR_1000022814)** — The switch may crash with a message similar to:
  `Software exception at alpha_chassis_slot_sm.c:563 -- in 'eChassMgr', task ID = 0`
- **Crash (PR_1000086916)** — The switch may crash with a message similar to:
  `Software exception at if_ether.c:693 -- in 'tNetTask', task ID = 0x196d9b0 -> ASSERT:  failed`
- **Crash (PR_1000087055)** — The switch may crash with a message similar to:
  `Software exception at ssh_utils.c:973 -- 'mftTask'`
- **Hang (PR_1000006985)** — The switch management may hang due to a memory corruption.
- **Security (PR_1000021329)** — Within the Web UI, the "Address Limit" value was always displayed as **4** for a learn mode of **Limited**.
- **Security (PR_1000021732)** — A configured IP Authorized Manager will fail following a reboot.
- **Security (PR_1000085928)** — The **show port-access authenticator 1** CLI command shows all port-access configured, but should show information for specified ports only.
- **SFlow (PR_1000021518)** — SFlow returns **sysUpTime** in 100ths of seconds, rather than 1000ths.

- **SFlow (PR_1000021776)** — The SFlow **sysUpTime** is not in sync with the switch **sysUpTime**.
- **SSH (PR_1000087086)** — The switch does not report an error message after rejecting a public key file with more than 10 keys.
- **Web UI (PR_89899)** — The Web UI port statistic counters are overwriting one another.
- **Web UI (PR_1000021867)** — VLAN context within Web UI may not allow untagged ports to be added to a VLAN.
- **Web UI (PR_1000085927)** — The Help text is not available from the authorized manager screen.

Release E.08.30

Problems Resolved in Release E.08.30

- **Auto-TFTP (PR_20802)** — Configuring the **auto-tftp** command with an incorrect IP address for the TFTP server can cause the switch to reboot every 5-15 minutes.
- **CLI (PR_1000000769)** — **update** and **upgrade-software** should not be normal CLI commands.
- **CLI (PR_1000001384)** — Misspelling in CLI Help screen for the **static-mac** command.
- **CLI (PR_1000001897)** — Help screen for **logging** command does not mention 'Major' logs.
- **CLI (PR_1000001628)** — The CLI may incorrectly reject the adding of ports to a VLAN, and respond with an `Inconsistent Value` error message.
- **CLI (PR_1000005912)** — The slot/module identifications within the CLI are incorrect and show slots numerically, rather than alphabetically.
- **CLI (PR_1000097427)** — Extraneous columns in the **show authentication** command.
- **Config (PR_1000020659)** — ProCurve 24 port 10/100 POE module identified with part number J8151A rather than with the appropriate part number, J8161A.
- **CDP (PR_1000004099)** — CDP advertises the switch as being only a router when routing is enabled. Changes made so that the switch now advertises itself as both a router and a switch when routing is enabled.
- **Crash (PR_1000007319)** — The switch may crash with a message similar to:

  ```
  Software exception in ISR at dmaRx.c:830 -> No resources available
  ```

- **Crash (PR_1000019386)** — The switch may crash with a message similar to:

```
NMI occured: IP=0x00466f68 MSR:0x0000b032 LR:0x00000000
Task='eDMAEmg001' Task ID=0x1625f58 cr: 0x22000000 sp:0x01625eb0
xer:0x00000000
```

- **Enhancement (PR_1000020429)** — Added the **show chassis-version** CLI command.
- **Help (PR_1000000560)** — Within the CLI, the Port Security Help file does not reference the learn mode "Limited-Continuous".
- **Help (PR_1000013464)** — The **show mac-address** Help text is too long and exceeds the 80 character limitation.
- **Monitoring Port (PR_1000012218)** — Port monitoring a mesh port can cause mesh packets to be transmitted out the wrong port.
- **PIM (PR_1000004117)** — "Expiry Time" changed to "Expire Time" following the **show ip PIM neighbor lists** command within the CLI.
- **PIM (PR_1000004818)** — PIM may not go into a forwarding state when a new neighbor that doesn't support state refresh connects.
- **PIM (PR_1000005019)** — PIM will forward state refresh that is not from the assert winner.
- **PoE (PR_1000019004)** — Extraneous `Power Denied` messages have been eliminated when EPS power has been lost.
- **Port Security (PR_1000013075)** — A port with Port Security enabled may learn addresses beyond its configured limit, and require a reboot to clear.
- **SSH (PR_1000003227)** — Need a special case for the SSH protocol-version configuration parameter to provide compatibility when back-revving to pre-E.08.*xx* code.
- **SSH (PR_1000004993)** — Memory corruption in SSH function.
- **SSL (PR_1000012823)** — SSL code modifications.
- **VLANs (PR_1000006670)** — Protocol VLANs configured in the CLI may not show up in the VLAN menu config screen and report that the member ports are orphaned.

Release E.08.07

Problems Resolved in Release E.08.07

- **ACL (PR_1000006679)** — The configured ACL "Range" parameter may not function properly after a reboot.
- **CLI (PR_82086)** — The command **show mac <mac-address>** does not function.

- **CLI (PR_ 1000000560)** — The port security “Help” screen has been updated to include learn mode "Limited-Continuous".
- **CLI (PR_1000004025)** — After the switch is up for approximately 49 days, the “Up Time” from the **show system** command will not be accurate.
- **CLI (PR_1000095690)** — Error message improved when a user enters an Interface Name that is too long.
- **Crash (PR_1000004216)** — The switch may crash with a message similar to:

  ```
  Driver corrupted - Slave Bus Error: dmaTxPollPackets.c:724
  ```

- **Crash (PR_1000005210)** — The switch may crash with a message similar to:

  ```
  Exception in ISR at dmaRx.c:830
  ```

- **Crash (PR_1000005829)** — The switch may crash with a message similar to:

  ```
  Software exception at alphaHwRateLimits.c:84
  ```

- **Crash (PR_1000005902)** — In cases where a heartbeat failure may occur, the switch will provide more specific and informative crash information.
- **Crash (PR_1000006392)** — The switch may crash with a message similar to:

  ```
  Software exception at pmgr_util.c:1500 -- in 'mLACPCtrl'
  ```

- **Crash (PR_1000006427)** — The switch may crash with a message similar to:

  ```
  Software exception at lacp_util.c:1723 - in 'mLACPCtrl'
  ```

- **Crash (PR_1000006833)** — The switch may crash with a message similar to:

  ```
  Slave crash at AlphaSlaveLearn.c:1576
  ```

- **Crash (PR_1000006967)** — The switch may crash with a message similar to:

  ```
  Exception at sw_malloc.c:141  Out of Memory - SSH
  ```

- **Crash (PR_1000006988)** — The switch may crash with a message similar to:

  ```
  Slave crash in ISR @ dmaRx.c:838
  ```

- **Crash (PR_1000007148)** — The switch may crash with a message similar to:

  ```
  Bus error:  HW Addr=0x7c7343b2 IP=0x002c3e54 Task='mIpAdMCtrl'
  ```

- **Crash (PR_1000007221)** — The switch may crash with a message similar to:

  ```
  Slave crash in mPmSlvCtrl at nc:phy.c:594
  ```

- **Crash (PR_1000007227)** — The switch may crash with a message similar to:

```
Software exception at alloc_free.c:485 -- in 'tDevPollTx', task ID =
0x17a3c58'
```

- **Crash (PR_1000011477)** — The switch may crash with a message similar to:

  ```
  Bus error:  HW Addr=0x06836252 IP=0x00444f14 Task='mHttpCtrl'  Task
  ID=0x11257f8
  ```

- **Crash (PR_1000011517)** — The switch may crash with a message similar to:

  ```
  Slave crash in ISR at dmaRx.c:838
  ```

- **Crash (PR_1000013156)** — Addressed master crash problem in memory system.
- **IP Helper (PR_1000004029)** — Number of IP Helper addresses increased to 256 on the 5300. See "Release E.08.42 Enhancements" on page 77.
- **MAC Authentication (PR_1000019250)** — The switch will crash if a MAC Authentication configured port is then configured for Trunking.
- **Meshing (PR_1000012101)** — A meshed switch may cause a broadcast loop on the network after a new module is inserted.
- **MDI/MDI-X (PR_1000001452)** - MDI/MDIX mode not described in help.
- **Port Monitoring (PR_1000012218)** — When port monitoring is configured, meshing protocol packets may be sent out the wrong meshed ports.
- **Port Security (PR_10000001437)** — Eavesdrop prevention. See "Release E.08.42 Enhancements" on page 77.
- **RMON (PR_1000011690)** — When RMON thresholds in the switch are exceeded, no trap is generated.
- **STP (PR_1000005371)** — Unable to set spanning tree "hello-time" via CLI in STP mode.
- **VLAN (PR_1000006670)** — If a port resides only in a protocol VLAN, the menu will not allow the user so save changes from the VLAN configuration window within the menu.
- **Web UI (PR_1000000256)** — The Web UI may display a module as a "humpback module"
- **Web UI (PR_1000007144)** — The VLAN Configuration help link is not available within the Web UI.

Release E.08.03

Problems Resolved in Release E.08.03

- **Crash (PR_1000007148)** — The switch may crash with a message similar to:

```
Bus error: HW Addr=0x7c7343b2 IP=0x002c3e54 Task='mIpAdMCtrl'.
```

- **Crash (PR_1000007227)** — The switch may crash with a message similar to:

```
Software exception at alloc_free.c:485 -- in 'tDevPollTx', task ID =
0x17a3.
```

Release E.08.01

Problems Resolved in Release E.08.01

- **ACL (PR_94945)** — 5300 allows duplicate ACEs (Access Control Entries) to be entered within an ACL.
- **CERT (PR_96648)** — Applied OpenSSH patches to switch for CERT Advisory CA-2003-24 related problems.
- **CLI (PR_81948)** — There are currently two “enable” commands present within the “Interface Config” context; one is to enable the port, the other is to enter manager context. The “enable” command is now filtered when not in the Operator Context within CLI.
- **CLI (PR_82475)** — The help text displayed for source-route is incorrect when auto-extend is applied to the command “IP”.
- **CLI (PR_90302)** — The help text within CLI for the “Interfaces” command is grammatically incorrect.
- **Crash (PR_88831)** — The switch may crash with a message similar to:

```
02/27/03 15:48:09 Bus error: HW Addr=0x02000000 IP=0x0013866c
Task='mSess2' Task ID=0x1654700fp: 0x01654a40 sp:0x016533a0
lr:0x0013874
```

- **Crash (PR_1000002979)** — The switch may crash with a message similar to:

```
Software exception at rstp_port_role_sm.c:44 -- in 'mRstpCtrl', task
ID = 0x1379a48-> ASSERT:  failed
```

- **Crash (PR_1000003288)** — The 10/100 Module (J4820A) under conditions of heavy port toggling may crash with a message similar to:

```
Software exception @ dmaRx.c: 237.
```

- **Crash (PR_89847)** — The switch may crash with a message similar to:

```
Software exception in ISR at alpha_hs_int.c:547

-> NCI_INTERRUPT_ERROR. Slot 3 NCI_IntReg=0x4000
```

- **Enhancement (PR_81844)** — Enhancement to improve “Debug HELP” information provided via CLI.
- **IP (PR_1000000728)** — The switch does not notify the IP Address Manager when an RSTP topology change occurs.
- **Logging (PR_82509)** — The switch will reboot when an invalid IP address is assigned to the logging feature, while “Logging” is turned off.
- **Meshing (PR_82502)** — Improved meshing performance during network conditions when there are large volumes of Port “learns” and “moves”.
- **Routing (PR_93205)** — The switch incorrectly allows for a configuration in which a static route can be configured as 127.x.x.x.
- **RSTP (PR_1000001612)** — A port takes approximately 30 seconds to go into the Forwarding state.
- **Security (PR_90899)** — After configuring a port to be "learn-mode configured", the "show port security" output within the CLI lists "Static" as the learn mode, rather than “Configured”, as it should be.
- **Security (PR_91855)** — The switch may fail to forward authentication requests to a RADIUS server when an unauthorized VID is configured and “Port-Security” is set to 802.1x.
- **Syslog (PR_91123)** — The switch may fail to send messages to a configured Syslog server.
- **VLAN (PR_92426)** — Unable to delete a VLAN by name if the name is numeric within the CLI.
- **Web Agent (PR_82157)** — There is a missing graphic in the upper left hand corner of the “First time installation” pop up window.
- **Web UI (PR_90858)** — Unable to clear the “VLAN Name” text field after 12 characters are entered within the Web UI.
- **XRMON (PR_98199)** — The “BroadcastPackets” counters for MIB object 1.3.6.1.2.1.16.1.1.1.6 on the 53xx series switch are incorrect.

Release E.07.40

Problems Resolved in Release E.07.40

- **Agent Hang (PR_97705)** — Agent processes (Ping, TELNET, SNMP, etc.) may stop functioning.
- **ARP (PR_92421/93008/97993)** — Default ARP aging time is 1,200 minutes when it should be 20 minutes. User-configured ARP aging times do work correctly.

- **Crash (PR_95293)** — The switch may crash with a message similar to:

  ```
  Bus error:  HW Addr=0x08000001 IP=0x00267cc4 Task='mIpAdMCtrl'  Task
  ID=0x150520 fp: 0x00000020 ip:0x01505100 lr:0x00267ca0
  ```

  This crash has been associated with traffic patterns generated by the Blaster and Welchia worms.

- **Crash (PR_96236)** — The switch may crash with a message similar to:

  ```
  "Software exception at ipaddrmgrSCtrl.c:2108 -- in 'mIpAdMUpCt'"
  ```

- **Crash (PR_97048/97083)** — The switch may crash with a message similar to:

  ```
  Bus error:  HW Addr=0x1bee13a8 IP=0x00267b68 Task='mIpAdMCtrl'  Task
  ID=0x14c2fe0 fp: 0x00000028 sp:0x014c2e98 lr:0x00267b58.  In QA code:
  Software exception at route.c:296. Attempt to free a null route.
  ```

- **Hang (PR_97031)** — Switch may hang (routing and console) due to infinite loop issue in ACL code.

- **Routing (PR_98494/97301)** — The switch may exhibit slower-than-normal routing performance due to route entries not being aged properly.

- **Routing (PR_98847)** — Under some conditions when there are more than 32 VLANs and IGMP enabled, the switch may not route.

- **X-modem (PR_95748)** — When trying to download a zero-length OS file to the switch, the switch may crash with a message similar to:

  ```
  Software exception at fileTransfer.c:552 -- in 'mftTask', task ID =
  0x1241ca8 -> Could not open file.
  ```

Release E.07.37

Problems Resolved in Release E.07.37

- **Crash (PR_90217)** — The switch may crash under high stress in a very large mesh topology with a message similar to:

  ```
  Bus error:  HW Addr=0x08040010 IP=0x002c8b48 Task='eDrvPoll' Task
  ID=0x177fdb0 fp: 0x01682e38 sp:0x0177f9e8 lr:0x002c8ae0.
  ```

- **Crash (PR_90374)** — The J4878A mini-GBIC module may cause the switch to crash with a message similar to:

  ```
  "Slot B SubSystem 0 went down:  01/01/90 13:05:41 Software exception
  at dmaRx.c:211 -- in 'tDevPollRx', task ID = 0x40808b78 -> FAULTY INK
  PARTNER CONNECTED ON SLOT".
  ```

- **Crash (PR_94852)** — The switch may crash when in a mesh configuration with a message similar to:

  ```
  Bus error:  HW Addr=0xdc37e837 IP=0x002c944c Task='eDrvPoll'  Task
  ID=0x173fdb0 fp: 0x01054468 sp:0x0173fa50 lr:0x002c93c0.
  ```

- **Crash (PR_95284)** — If a user enters an invalid MAC address during the Port Security configuration within the CLI, the switch may crash with a message similar to:

  ```
  Software exception at exception.c:345 -- in 'mSess1', task ID = 0x141ae70
  -> Memory system error at 0x131b5a0 - memPartFree
  ```

- **Meshing (PR_96007)** — If a mesh link is broken then shortly followed by the learning of new MAC addresses, the switch may exhibit problems such as bus errors and/or improper communication with other mesh switches.

- **SNMP (PR_96999)** — When the switch is reset (or power-cycled) after configuring an SNMP Community Name with “Operator/Restricted” Rights, it will still allow SNMP sets (writes) to MIB objects.

Release E.07.34

Problems Resolved in Release E.07.34

- **Agent Hang (PR_92802)** — The switch may become unresponsive or hang due to UDP port 1024 broadcast packets never being freed, after the TIMEP and SNTP clients are disabled on the switch.

- **Crash (PR_92659)** — Software exception at memrpt.c:1153 -- in 'mInstCtrl', task ID = 0x1455a30

- **IPv6 (PR_93171)** — The switch does not forward IPv6 Router Solicitation/Advertisements when IGMP is enabled.

- **Routing/Agent Performance (PR_95009)** — Routing performance may be degraded due to the aging of host route entries. In this scenario, traffic will be routed through the switch software, thereby resulting in lower performance of routing and agent access (TELNET, SNMP, ping, etc.) operations.

- **VLAN (PR_92466)** — The switch may experience a Bus error related to 802.1X/unauthorized VLAN. The Bus error is similar to:

  ```
  Bus error:  HW Addr=0x3861000c IP=0x002df470 Task='mAdMgrCtrl'Task
  ID=0x16e616 0 fp: 0x006a090c sp:0x016e5df0 lr:0x0021d6d8
  ```

Release E.07.30

Problems Resolved in Release E.07.30

- **Agent Performance (PR_81861)** — The switch may get into a state where end nodes and other network devices cannot contact (ping, telnet, SNMP, etc.) the switch's agent.
- **Routing (PR_90802/91236)** — The switch may route packets out the wrong port due to a mismatch between the source and destination MAC addresses.

Release E.07.29

Problems Resolved in Release E.07.29

- **ACL (No PR)** — The switch allows a user to execute a “no access-list” command for a non-existent ACE without responding with an appropriate error message.
- **ACL (PR_90250)** — Packets that match a “denied” ACL entry may cause the switch’s CPU to run at full utilization.
- **ACL (PR_90415)** — On ACL entries such as “permit/deny tcp any any” the switch will incorrectly permit/deny UDP traffic. The same is inversely true for ACL entries such as “permit/deny udp any any” resulting in TCP traffic being permitted/denied.
- **ACL Performance (PR_90366)** — Addressed potential performance issues of cached TCP and UDP ACL entries.
- **Config (PR_88753)** — A 1000-FDX port setting in the switch config file is not processed properly, resulting in Gigabit-SX ports remaining in an “auto” port configuration. This is most often seen when reloading or TFTP’ing a config file to the switch.
- **Port-sec (PR_88612)** — Static MAC addresses are set up under port security with learn-mode “configure specific” to allow those MAC addresses to communicate through the switch. If one of those MAC addresses is removed via the Web interface of the switch and then re-entered, the owner of that MAC address cannot communicate through the switch.
- **Routing (PR_90554)** — Cached routing information was only updated by IP routable datagrams, and was not being updated by Layer-2 traffic such as ARP.
- **Self Test (PR_90777)** — A self test error may occur when a Gigabit-SX, or LX mini-GBIC module is inserted into the switch while powered on.
- **Spanning Tree (PR_90412)** — Enhancements made to 802.1w operation to address version 3 BPDU communication issues.

Release E.07.27

Problems Resolved in Release E.07.27

- **Enhancement (PR_90365) —** Modifications have been made to the switch meshing code to allow limited mesh interoperability between E.07.*x* and pre-E.07.*x* software to allow easier upgrades of all switches in a mesh. (Current implementation does not allow switches running pre-E.07.*x* software to participate with meshed switches running E.07.*x* or greater.)
- **IGMP (PR_82491) —** A Group-Specific Query (GSQ) timeout is currently .2 to .6 seconds, rather than the specified default of 1 second.
- **IGMP (PR_90376) —** In some cases, the switch would display “0.0.0.0” for the output of the CLI command “show ip igmp”.
- **Meshing (PR_88689) —** A 12-switch mesh may cause the switch to temporarily run out of packet buffers.
- **Telnet (PR_82522) —** Switch TELNET connections were not closed properly resulting in new TELNET sessions being established which could result in the switch reaching its maximum number (3) of TELNET sessions.
- **Web-Browser Interface (PR_82530) —** A client using Sun java 1.3.X or 1.4.X to access the Web-Browser Interface of the switch, may cause the switch's CPU utilization to increase causing agent processes (such as console, telnet, STP, ping, etc.) to stop functioning.

Release E.07.22

Problems Resolved in Release E.07.22

- **Meshing/Packet Buffer Depletion (PR_88694) —** Certain mesh topologies may cause packet buffers to be depleted on the switch. In this state the switch will generate an "`Out of pkt buffers`" Event Log message.
- **OSPF (PR_88718) —** In topologies where the switch has redundant routes (via a directly connected link and via an OSPF learned route) to the same network, the switch does not learn the alternate route via OSPF when the directly connected link goes down.
- **Port Hang (Packet Not Forwarded) (PR_88613) —** Under certain traffic load conditions, ports that are toggling on the mini-GBIC module (J4878A) may stop transmitting packets.

Release E.07.21

Problems Resolved in Release E.07.21

- **ARP (PR_5185) —** ARP has been enhanced to have a configurable timeout value, beyond the current default of 20 minutes.
- **CDP (PR_5054) —** CDP multicasts are not passed when CDP is disabled on the switch.
- **CLI (PR_5053) —** Setting the telnet inactivity timeout from the CLI does not indicate a reboot is necessary for changes to take effect.
- **CLI (PR_4984) —** The definition of default gateway following the “ip?” in the CLI is stated as “Add/delete default route to/from routing tale.”, which is incorrect. Clarified help text for 'ip default-gateway' CLI command to state that this parameter is only used if routing is not enabled on the switch.
- **CLI (PR_5242) —** Information in the command “show boot-history” is not in the order claimed (most recent first).
- **Crash (PR_4621) —** The switch may crash with a message similar to:

  ```
  NMI occurred: IP=0x00317d9c MSR:0x0000b000 LR:0x00013b88
  Task='eDrvPollRx' Task ID=0x1708f20 cr: 0x22000080 sp:0x01708e60 xer:
  ```

- **Crash (PR_5745) —** The switch may crash with a message similar to:

  ```
  -> Divide by Zero Error: IP=0x801400c0 Task='sal_dpc_hi'
     Task ID=0x80616690 fp:0x00000000 sp:0x80616600 ra:0x800140060
     sr:0x1000af01
  ```

- **Crash (PR_5635) —** The switch may crash with a message similar to:

  ```
  -> Assertion failed:0, file drvmem.c,  line 167
  ```

- **Crash (PR_5679) —** The switch may crash with a message similar to:

  ```
  -> Bus error: HW Addr=0x00000000 IP=0x00000000 Task='mNSR' Task
     ID=0x1725148 fp: 0x0000c4b0 sp:0x012e9780 lr:0x00330674
  ```

- **Crash (PR_5712) —** The switch may crash with a message similar to:

  ```
  -> TLB Miss: Virtual Addr=0x00000000 IP=0x8002432c Task='tSmeDebug'
  ```

- **Crash (PR_5725) —** The switch may crash with a message similar to:

  ```
  -> Assertion failed: nt, file dpc.c, line 169
  ```

- **Crash (PR_5846) —** WhatsUpGold telnet scan can cause switch to run out of memory and crash with error message similar to:

  ```
  -> malloc_else_fatal() ran out of memory
  ```

- **Crash (PR_5955) —** The switch may crash with a message similar to:

  ```
  Software exception at alpha_chassis_slot_sm.c:506
  ```

- **Crash (PR_4986) —** The switch may crash with a message similar to:

  ```
  -> Bus error: HW Addr=0x00ffffff IP=0x332c4530 Task='mSess1' Task
     ID=0x16a62f0 fp: 0x2e2e2e29 sp:0x016a61a0 lr:0x0010f028
  ```

  This crash can occur when eight transceiver modules are installed and the command “interface all” is typed in the configuration context.

- **Crash (PR_5418) —** The switch may crash with a message similar to:

  ```
  -> Software exception at rtsock.c:459 -- in 'tNetTask', task ID =
  0x1a225b0
  ```

- **Crash (PR_5635) —** The switch may crash with a message similar to:

  ```
  -> Assertion failed:0, file drvmem.c,  line 167
  ```

- **Crash (PR_5341) —** All three of the following steps must occur before the crash is exhibited:

  1. .A 1000-T port (without a link) is configured as a mirror destination port.
  2. Another blade/port traffic is mirrored to that destination port.
  3. Mirror destination port/blade will crash or hang after connecting, then disconnecting a 100T link with a message similar to:

     ```
     Software exception at nc_fd_fi.c:693 - in 'mPmSlvCtrl'task ID =
     0x405e9cc8 -> netchip_FIOutboundFlush: Timeout reached!
     ```

- **Crash (PR_5236) —** The switch may crash with a message similar to:

  ```
  -> AlphaSlaveAddrmgr.p 1021 this time
  ```

  This crash can occur when a module is hot-swapped after downloading new software to the switch without rebooting.

- **Date/Time (PR_5264) —** The timezone can cause the date to wrap if the timezone is set to a valid, but negative value (like -720) without previously configuring the switch's time. The switch may report an invalid year (i.e. 2126).

- **DHCP —** If a client moves without first releasing its IP address, it will not receive a NAK, resulting in the client’s inability to get an IP address at its new location.

- **Event Log (PR_5154) —** When a module fails to download, the severity code is INFO instead of WARNING.

- **Fault Finder/CLI (PR_4696) —** Setting fault finder sensitivity always resets action configuration to 'warn', when it should remain 'warn and disable'.

- **FFI/Port Counters (PR_5429) —** No errors are reported by the FFI or port counters when linking at 100 HDX on a Gigabit port with a duplex mismatch.
- **FFI/Port counters (PR_5280) —** FFI and port counters don't have consistent values.
- **Filter (PR_5132) —** Source port Filter on Dyn1 LACP trunk creates Multicast Filter entry that cannot be deleted.
- **Filter (PR_4833) —** Creating a source port filter for a port, moving the port into a trunk, and then reloading the saved TFTP configuration file results in a corrupted download file error.
- **Flow Control (PR_5102) —** Setting a port “X1” in 10-HDX, then attempting to turn on flow control returns an error similar to: “Error setting value fl for port X2”. The error should read “X1”.
- **GVRP (PR_5284) —** Port does not register VLAN even though advertisements are received.
- **Hot-swap (PR_4900) —** Hot-swapping a transceiver logs a message requesting to reboot the switch in order to enable the port, which is not necessary.
- **IGMP (PR_5736) —** If IGMP is turned on for multiple VLANs, and is then turned off for a single VLAN, the Data-Driven Mcast filters for that VLAN are not flushed.
- **IP (PR_5408) —** IP is causing the driver to apply source port filters incorrectly to non-routed packets.
- **IRDP (PR_5923) —** When running the 'rdisc' router discovery tool under Redhat 8.0 or 7.3, Linux reports “ICMP Router Advertise from <IP>: Too short 16 40” when a IRDP packet is received.
- **LACP/Port Security (PR_5059) —** With LACP on, the command “port-sec a1 l c action send-alarm” fails with a message similar to “learn-mode: Inconsistent value”.
- **Link Toggle Corruption (PR_5527) —** Addressed issue whereby toggling ports with active, bi-directional traffic could result in corrupted packets within the system.
- **Link-up Polling Interval (PR_5000) —** A delay of up to 1.7 seconds between plugging in a cable (linkbeat established) and traffic being forwarded to and from that port may cause problems with some time sensitive applications. For example, AppleTalk dynamic address negotiation can be affected, resulting in multiple devices using the same AppleTalk address.
- **Menu (PR_5346) —** The one-line help text below the password entry field, displays the message "`Enter up to 16 characters (case sensative), or just press <Enter> to quit`". It should read "`...sensitive...`".
- **Meshing (PR_4969) —** Traffic on oversubscribed mesh links will migrate to other mesh links too slowly.

- **Meshing (PR_4980) —** Meshing does not maintain priority on encapsulated packets that are sent out non-mesh ports.
- **Multicast Filters (PR_4741) —** Any static multicast filters configured once the limit has been reached, would appear in the output of the "show filter" CLI command with only partial information. Switch now correctly returns error message "Unable to add filter" once limit has been reached.
- **OSPF (PR_88611) —** When configured for authentication-key type "simple passwords", the switch does not include the password in OSPF packets.
- **Port Configuration (PR_5444) —** When interchanging 10/100-TX modules J4862A and J4862B, the port configuration of the module originally installed in the switch is lost.
- **Port counters (PR_5013) —** Hardware port counter filters for dot1dTpPortInDiscards not implemented.
- **Port counters (PR_5171) —** The "Total RX Error" counter is incorrect when the port has heavy 10HDx traffic.
- **Port counters (PR_5204) —** The Runt Rx counter in the detail port counter screen, does not increment when there are fragments.
- **Port counters (PR_5400) —** The 64-bit counter for the highest numbered port on a given module, does not update properly.
- **RADIUS (PR_4886) —** Pressing the tab key gives error message similar to "BAD CHARACTER IN ttyio_line: 0x9n" when entering a username for the radius prompt.
- **RSTP (PR_5449) —** There is a delay in the switch relearning MAC addresses when an RSTP port transitions from Blocking to Forwarding.
- **Self Test (PR_5113) —** There are intermittent port failures reported on ProCurve switch xl 100/1000-T modules (J4821A) while performing a packet self test, which was due to the packet test not seeing the very first packet.
- **SNMP (PR_5349) —** The switch does not send SNMP packets larger than 484 bytes.
- **SNTP/TIMEP (PR_5018) —** SNTP still runs when TIMEP is enabled.
- **Source Port Filters (PR_4669) —** Source port filters for illegal ports and trunk port members cannot be deleted from the CLI.
- **Source Port Filters (PR_4719) —** The switch does not automatically remove a source port filter for a trunk that has been deleted.
- **System Information (PR_5169) —** Up Time displayed is not correct.
- **TACACS (PR_5226) —** During TACACS Authentication the TACACS Server's IP address is shown on the switch's 'splash screen'.

- **TCP (PR_5227) —** TCP port 1506 is always open. Fix is to close TCP port 1506.
- **TFTP (PR_5034) —** Trying to TFTP a config onto the switch causes the switch to not complete its reload process. The switch hangs and does not come up.
- **VLANs (PR_4405) —** The VIDs of deleted VLANs are not removed from the switch's VLAN table, causing the switch to not allow new VLANs to be created (once the VID table is full).
- **Web (PR_5455) —** Bad URL was being mirrored back to the user following Nessus script attack test.
- **Web-Browser Interface (PR_5199) —** Having a ProCurve switch 4100gl series as a commander, and a ProCurve switch 4000m as a member of the stack, the stack commander was not checking security when doing passthrough.
- **Web-Browser Interface (PR_5052) —** The CLI does not disable the web-browser interface.
- **Web-Browser Interface (PR_5055) —** Missing firmware/ROM information in Web UI.
- **Web-Browser Interface (PR_5158) —** When clicking on the Web UI System Info "Apply Changes" button, a character appears under the "VLAN Configuration" tab.
- **Web-Browser Interface (PR_4976) —** Mis-spelled word on the product registration screen of the WEB UI. The phrase "...does not appears above..." is now "...does not appear above..."
- **Web-Browser Interface (PR_4996) —** When using a ProCurve Switch 4108 as a commander switch in the stack, a ProCurve Switch 2424M is not shown in the device view of the stack closeup in the web UI. The message "Device view, HP2424M, not supported by firmware of commander" is present instead of the device view.
- **Web-Browser Interface (PR_4904) —** When a transceiver is removed from the switch, its configuration is not cleared on the Status->port status screen of the web UI. The transceiver type will still show until a new transceiver is inserted.
- **Web-browser Interface (PR_4235) —** Web-browser port utilization label does not display the bandwidth number. Shows x% of 0Mb instead of x% of 100Mb or x% of 1Gb.
- **Web-Browser Interface (PR_4495 —** Administrator password can be used in combination with the operator username.

Release E.06.10

Problems Resolved in Release E.06.10

- **Crash (PR_5229) —** Greater than 100 hotswaps causes mesg buff crash.

- **Flow Control (PR_5215) —** Enabling Flow Control on a port does not enable Global Flow Control on the switch.
- **Security (PR_5226) —** Removed display of TACACS Server IP address during remote management logon.
- **Security (PR_5227) —** TCP Port 1506 access is closed when Telnet or Stacking is disabled.
- **Web-browser interface (PR_5052) —** Executing the CLI command "**no web-management**" does not disable access to the web-browser interface.

Release E.06.05

Problems Resolved in Release E.06.05

- **Crash (PR_5471)** — The CLI command "show ip ospf neighbor" may cause the switch to crash with a message similar to:

```
Bus error: HW Addr=0x30008fa0 IP=0x001112a4 Task='mSess1' Task
ID=0x169b110
```

Release E.06.03

Problems Resolved in Release E.06.03

- **Packets not Forwarded (PR_5201)** — A synchronization issue between the switch chassis and modules after several weeks of continuous operation can result in packets being dropped by the switch instead of being forwarded.

Release E.06.02

Problems Resolved in Release E.06.02

- **Performance (PR_5161)** — Certain high traffic levels may cause the switch to drop packets.

Release E.06.01

Problems Resolved in Release E.06.01

- **100/1000-T module (PR_4956) —** Bringing a port up and down while the port is running at or near maximum throughput may cause the module to reset.
- **802.1x (PR_4972) —** Support for 802.1x is not implemented in routing mode.
- **802.1x (PR_5043) —** When changing an 802.1x port configuration, the switch does not correctly restore default VLAN ID after disconnecting the port.
- **ARP (PR_4443) —** Switch incorrectly replied to an ARP packet with a header length ranging from 7 to 15 bytes. The switch now replies only if header length is equal to 6 bytes.

- **CDP (PR_4546) —** CDP multicast packets are not passed through the switch when CDP is disabled on the switch.
- **CLI/RIP (PR_5046) —** The CLI command 'show ip rip interface' results in the following:

  ```
  "RIP interface information for 0.0.0.0, RIP is not configured on this..."
  ```

- **CoS (PR_4738) —** Cannot configure CoS on a trunk port. Also, enhancements to CoS error handling when moving ports in and out of a trunk.
- **CoS (PR_4982) —** The output of the CLI command “show qos port-priority” may show an illegal state (“no priority”) for the Differentiated Services Codepoint (DSCP) policy. This problem may occur given this situation:

  1. Configure a DSCP policy on a port, and
  2. Remove module, and
  3. Reboot switch, and
  4. Delete DSCP policy, and
  5. Hot-swap module back into the switch

- **Crash (PR_4933) —** Switch may crash while hot swapping a module with a message similar to:

  ```
  -> Software exception in ISR@alloc_free.c:479
  ```

- **DHCP-Relay (PR_4551) —** Configuring an IP helper address on a VLAN does not automatically turn on the DHCP-relay function.
- **Extended RMON (PR_5083) —** When Extended RMON and Routing are enabled, the switch may duplicate packets on the network.
- **LACP (PR_5000) —** Link-up polling interval: A delay of up to 1.7 seconds between plugging in a cable (linkbeat established) and traffic being forwarded to and from that port may cause problems with some time sensitive applications. For example, AppleTalk dynamic address negotiation can be affected, resulting in multiple devices using the same AppleTalk address.
- **Mini-GBIC Link Connectivity Issue (PR_4957) —** A mini-GBIC Gigabit-SX/LX link between an ProCurve Switch 5300xl and an ProCurve Routing Switch 9300 may not be established when both sides are in the default configuration (Auto).
- **Radius (PR_4886) —** If using the TAB key while entering a username for the radius prompt, the switch may display an error message similar to:

  ```
  ->BAD CHARACTER IN ttyio_line: 0x9n
  ```

- **RIP (PR_4757) —** After the switch reboots and if a routing loop (3 or more routers) exists in the topology, RIP may age out its own connected routes (even though the routes are still valid).

- **RIP (PR_4965) —** Static routes are redistributed into RIP. [Fix: Static routes are no longer redistributed into RIP by default, only directly connected routes are redistributed.] [Old description: Changes to RIP route redistribution such that only connected routes are redistributed, not static configured routes.
- **RIP (PR_4987) —** If multiple IP addresses are configured for a VLAN and RIP is running on one or more of the secondary addresses, the CLI command “show ip rip vlan x” will only show information about the primary IP address.
- **Routing (PR_4977) —** If a default route is not configured and the switch receives a Layer 3 packet with an unknown source address, the packet will be routed by software even though an entry for the destination exists in the hardware routing table.
- **Static Routes (PR_5040) —** Reject static routes could not be created.
- **Web Browser Interface (PR_4976) —** The product Registration screen contains a typographical error. The phrase “...does not appears above...” is now “...does not appear above...”.

Part Number 5991-2127
January, 2006