



Release Notes: Version I.08.62 Software *for the ProCurve 2800 Series*

Release I.08.62 supports these switches:

- ProCurve Switch 2824 (J4903A)
- ProCurve Switch 2848 (J4904A)

These release notes include information on the following:

- Downloading switch documentation and software from the Web ([page 1](#))
 - Clarification of operating details for certain software features ([page 8](#))
 - Software features available in release I.07.*xxx* and I.08.*xxx* ([page 11](#))
 - A listing of software fixes included in releases I.07.31 through I.08.62 ([page 30](#))
-

Boot ROM Update Required

A successful update to I.08.*xx* requires updating the 2800 with the current Boot ROM version, I.08.05. The I.08.05 Boot ROM image was automatically installed on any switch running the I.07.68 software. If your 2800 is currently running a pre-I.07.68 software version, you must update the Boot ROM before installing I.08.*xx*. Load the I.07.68 software and reboot your switch from the I.07.68 software image. NOTE: a copy of the I.07.68 software is included in the I.08.*xx* zip file on the ProCurve Networking Web site. See "[Release I.07.64 Enhancements](#)" on [page 26](#) for more information.

Caution

The startup-config file saved under version I.08.*xx* or greater, is NOT backward-compatible with previous software versions. Users are advised to save a copy of the pre-I.08.*xx* startup-config file BEFORE UPGRADING to I.08.*xx* or greater, in case there is ever a need to revert to pre-I.08.*xx* software. Instructions for saving a copy of the startup-config file are found in the "Transferring Switch Configurations" section of Appendix A in the *Management and Configuration Guide* (included in PDF format on the Product Documentation CD-ROM) shipped with the switch, and also available on the ProCurve Networking Web site. (Refer to "[To Download Product Documentation:](#)" on [page 1](#).)

© Copyright 2001, 2005 Hewlett-Packard Company, LP.
The information contained herein is subject to change
without notice.

Publication Number

5990-6049
April 2005

Applicable Product

ProCurve Switch 2824	(J4903A)
ProCurve Switch 2848	(J4904A)

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](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
<http://www.hp.com/go/procurve>

Contents

Software Management	1
Downloading Switch Documentation and Software from the Web	1
Downloading Software to the Switch	2
TFTP Download from a Server	3
Xmodem Download From a PC or Unix Workstation	4
Saving Configurations While Using the CLI	6
ProCurve Switch, Routing Switch, and Router Software Keys	7
Minimum Software Versions for Series 2800 Switch Features	7
Clarifications	8
IGMP	8
Supported Standards and RFCs	8
IGMP, Multicast Filters, and Configured VLANs	8
Using Delayed Group Flush	9
Setting Fast-Leave and Forced Fast-Leave from the CLI	9
IGMP Operating Notes	10
Displaying Spanning Tree Configuration Detail	10
Enhancements	11
Release I.08.62 Enhancements	11
Release I.08.61 Enhancements	11
Release I.08.60 Enhancements	11
RADIUS Authentication for Web Browser Access	12
CLI Local Terminal Mode	12
DHCP Option 82	13
Release I.08.55 Enhancements	25
Release I.07.68 Enhancements	26
Boot ROM Update	26
Release I.07.64 Enhancements	26
Boot ROM update	26
Release I.07.59 - I.07.63 Enhancements	27

Release I.07.58 Enhancements	27
Boot ROM update	27
Release I.07.53 - I.07.57 Enhancements	27
Release I.07.52 Enhancements (Beta Only)	28
QOS Pass-Through Mode	28
Release I.07.51 Enhancements	28
Release I.07.50 Enhancements	28
Port Trunking	28
Port Monitoring	29
Release I.07.32 Enhancements	29
Software Fixes in Release I.07.xx - I.08.xx	30
Release I.08.62	30
Release I.08.61	30
Release I.08.60	31
Release I.08.58	32
Release I.08.57	32
Release I.08.56	32
Release I.08.55	33
Release I.07.68	34
Release I.07.67	34
Release I.07.66	35
Release I.07.65 (Not Released)	35
Release I.07.64	35
Release I.07.63 (Beta Only)	35
Release I.07.62 (Beta Only)	35
Release I.07.61	36
Release I.07.60	36
Release I.07.59 (Beta Only)	36
Release I.07.58 (Beta Only)	36
Release I.07.57 (Never Released)	37
Release I.07.56	37

Release I.07.55 (Beta Only)	37
Release I.07.53 (Beta Only)	38
Release I.07.52 (Beta Only)	38
Release I.07.50	38
Release I.07.32	40
Known Software Issues and Limitations	42
Issues	42
Limitations	42
Displaying the Fast-Leave Setting on a Port	42

— *This page is intentionally unused.* —

Software Management


Downloading Switch Documentation and Software from the Web

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

To Download a Software Version:

1. Go to the ProCurve Networking Web site at:
<http://www.hp.com/go/procurve>.
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 Networking Web site at <http://www.hp.com/go/procurve>.
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 the “[Caution](#)” on the front page.

HP periodically provides switch software updates through the ProCurve Networking Web site (<http://www.hp.com/go/procurve>). After you acquire the new software file, you can use one of the following methods for downloading the software 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).
- HP’s SNMP Download Manager included in ProCurve Manager
- A switch-to-switch file transfer

Note

Downloading a new software version 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 I_08_62.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 I_08_62.swi
Device will be rebooted, do you want to continue [y/n]? y
02304K _
```

2. When the switch finishes downloading the software file from the server, it displays this progress message to indicate the switch is writing the downloaded software to flash memory:

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

3. After the switch writes the downloaded software to flash memory (approximately three minutes), you will see this screen:

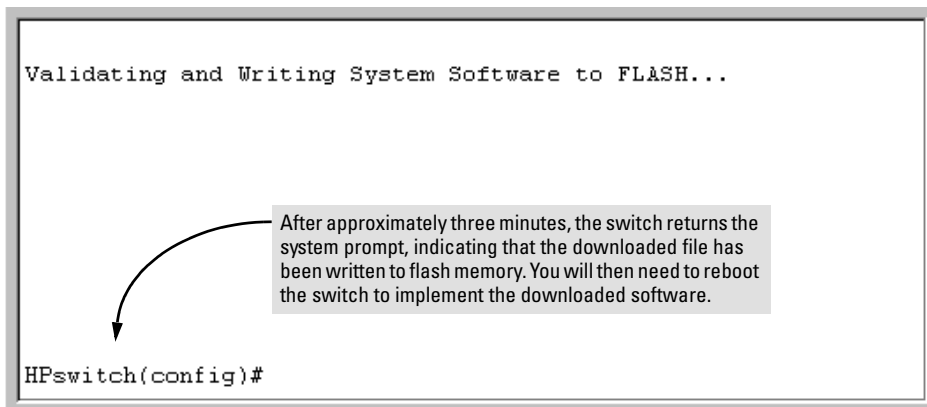


Figure 1. The System Prompt Indicates the Switch Is Ready To Activate the Downloaded Software

4. Type the **show flash** command to verify that the new software version has been downloaded to flash:

```
HPswitch# show flash
Image                Size(Bytes)  Date      Version
-----
Primary Image       : 2942618   02/08/05  I. 08. 62
Secondary Image     : 2435164   01/21/05  I. 07. 68
Boot Rom Version:   I. 08. 05
Current Boot       : Primary
```

← Software Version in Primary Flash

Figure 2. Example of Show Flash Display of Software Versions in Flash Memory

5. Use the CLI **boot** command to reboot the switch.

```
HPswitch(config)# boot
Device will be rebooted, do you want to continue [y/n]?
```

To continue the boot process, press the **[Y]** key (for “yes”). 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 can increase the baud rate to a value such as 57600 bits per second in your terminal emulator and in the switch. The baud rate must be the same in both devices. For example, to change the baud rate in the switch to 57600, execute the following commands:

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

Reboot your switch and re-establish a console session at the higher baud rate. Be sure to set your terminal emulator to the same baud rate.

2. Execute the following command in the CLI:

```
HPswitch(config)# copy xmodem flash pc
Device will be rebooted, do you want to 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.

When the download finishes, the switch automatically reboots itself and begins running the new software version.

4. To confirm that the software downloaded correctly:

HPswitch> show system

Check the **Firmware revision** line.

```
HPswitch# Show system
Status and Counters - General System Information

System Name       : HP ProCurve Switch 2824
System Contact    :
System Location   :

MAC Age Time (sec) : 300

Time Zone         : 0
Daylight Time Rule : None

Firmware revision : I.08.62      Base MAC Addr   : 00057-fdae40
ROM Version       : I.08.05      Serial Number    : SG00001588

Up Time           : 2 hours      Memory - Total  : 33,037,936
CPU Util (%)      : 3            Memory - Free   : 25,482,512

IP Mgmt - Pkts Rx : 334          Packet - Total  : 1998
           Pkts Tx : 69          Buffers - Free  : 1749
                                   Lowest  : 1738
                                   Missed  : 0
```

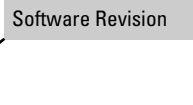


Figure 3. Example of Using the CLI 'show system' Command to Verify the Software Revision

5. 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.) (Remember 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, press **[Y]** (for Yes) when you see the “save configuration” prompt:

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

ProCurve Switch, Routing Switch, and Router Software Keys

Software Letter	ProCurve Switch, Routing Switch, or Router
C	1600M, 2400M, 2424M, 4000M, and 8000M
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 <ul style="list-style-type: none"> • H.07.50 and Earlier • H.08.55 and 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 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 2800 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.hp.com/go/procurve>.
2. Click on **software updates**.
3. Click on **Minimum Software Version Required by Feature**.

For Switch 2800 Hardware Accessories.

ProCurve Device	Minimum Supported Software Version
J4858A Gigabit-SX-LC Mini-GBIC	I.07.31
J4859A Gigabit-LX-LC Mini-GBIC	I.07.31
J4860A Gigabit-LH-LC Mini-GBIC	I.07.31
J8168A ProCurve 600 RPS/EPS	I.07.31

Clarifications

IGMP

Note: the following information updates and clarifies information in Chapter 4, “Multimedia Traffic Control with IP Multicast (IGMP)” in the *Advanced Traffic Management Guide*—part number 5990-8853, October 2004. Please review this chapter for a detailed explanation of IGMP operation.

Supported Standards and RFCs

The following are supported:

- RFC2236 (IGMP V.2, with backwards support for IGMP V.1)
- Interoperability with RFC3376 (IGMPv3)
- IETF draft for IGMP and MLD snooping switches (for IGMP V1, V2 V3)

The switch provides full IGMPv2 support as well as full support for IGMPv1 Joins. The switch is interoperable with IGMPv3 Joins as it forwards packets for the joined group from all sources. It does not support IGMPv3 “Exclude Source” or “Include Source” options in the Join Reports. The switch can operate in IGMPv2 Querier mode on VLANs with an IP address.

IGMP is supported in the HP MIB, rather than the standard IGMP MIBs, as the latter reduce Group Membership detail in switched environments.

IGMP, Multicast Filters, and Configured VLANs

On the Series 2800 Switches, the number of multicast filters available for use by IGMP is a function of the number of VLANs and the number of IGMP-enabled VLANs configured on the switch. When the number of multicast groups on the network exceeds the number of multicast filters available on the switch, excess multicast group traffic is flooded to all ports on the affected VLAN, to ensure that clients can receive their multicast traffic.

There are 255 multicast filters available in the switch. A filter is used each time a VLAN is configured and each time IGMP is enabled on a VLAN. The table below shows examples of the number of multicast filters available for IGMP use, based on the number of configured VLANs and IGMP-enabled VLANs.

Configured VLANs	IGMP-Enabled VLANs	Multicast Filters Available for IGMP	Average Number of Multicast Filters Available for IGMP per IGMP-Enabled VLAN
256	1	$255 - 256 - 1 = -2$	None. IGMP does not operate.
64	64	$255 - 64 - 64 = 127$	< 2
30	30	$255 - 30 - 30 = 195$	6.5
64	8	$255 - 64 - 8 = 183$	< 23
100	2	$255 - 100 - 2 = 153$	76

Using Delayed Group Flush

This feature continues to filter IGMP-Left groups for a specified additional period of time. This is beneficial in switches such as the Series 2600 or 4100gl, where Data-Driven IGMP is not supported. The delay in flushing the group filter prevents stale traffic from being forwarded by the server. Delayed Group Flush is enabled or disabled for the entire switch.

As the Series 2800 Switches use Data-Driven IGMP with IGMP Fast-Leave always enabled, HP does **not** recommend that the Delayed Group Flush feature be used on the Series 2800 Switches. Note that this command must be executed in the configuration context.

Syntax: `igmp delayedflush <time period>`

*Enables the switch to continue to flush IGMP-Left groups for a specified period of time (0 - 255 seconds). The default setting is **Disabled**. To disable, reset the time period to zero.*

Syntax: `Show igmp delayedflush`

Displays the current setting for the switch.

Setting Fast-Leave and Forced Fast-Leave from the CLI

In previous software versions, Fast-Leave and Forced Fast-Leave options for a port were set through the MIB. The following commands now allow a port to be configured for fast-leave or forced fast-leave operation from the CLI. Note that these command must be executed in a VLAN context

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

*Enables IGMP Fast-Leave on the specified ports in the VLAN (the default setting). In the Config context, use the VLAN specifier; for example, **vlan < vid > ip igmp fastleave <port-list>**. The “no” form disables Fast-Leave on the specified ports.*

[no] ip igmp forcedfastleave <port-list>

Forces IGMP Fast-Leave on the specified ports in the VLAN, even if they are cascaded. The “no” form disables Forced Fast-Leave on the specified ports.

To view the IGMP Forced Fast-Leave status of a port use the **show running-config** or **show configuration** commands.

IGMP Operating Notes

- Review the number of VLANs and the number of IGMP-enabled VLANs you plan to use to determine if you have the sufficient multicast filters available for your expected IGMP groups. If you don't, excess multicast groups are not filtered and are flooded to all ports on the VLAN.
- Do not use Delayed Group Flush for Series 2800 Switches, as this behavior provides no additional benefits when Data-Driven IGMP is supported.
- Forced fast leave can be used when there are multiple devices attached to a port.

Displaying Spanning Tree Configuration Detail

A new CLI command has been added to provide more detailed statistics on spanning tree operation.

Syntax: show spanning-tree <port-list> detail

Lists 802.1D and 802.1w port operating statistics for all ports, or those specified.

Enhancements

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

Release I.08.62 Enhancements

I.08.62 Enhancement	Overview
QoS Pass-Through can be enabled dynamically.	This feature can now be configured without requiring a switch reboot to activate. For more information on QoS Pass-Through, see page 28 . An updated release note with more information will be provided in March, 2005.

Release I.08.61 Enhancements

Software fixes only; no new enhancements.

Release I.08.60 Enhancements

I.08.60 Enhancement	Overview
RADIUS Authentication for Switch 2800 Web Browser Access	The aaa authentication command now allows the optional use of RADIUS as the primary password authentication method for the Web browser interface on Series 2800 switches (as well as for the Series 2600, 2600-PWR, and 5300xl switches). Refer to " RADIUS Authentication for Web Browser Access ", below
CLI Local Terminal Mode Command	This new command enables changing from one terminal mode to another without changing the terminal mode configuration or having to reboot the switch. The command is not persistent across reboots, and affects only the current console session. Refer to " CLI Local Terminal Mode " on page 12 .
DHCP Option 82	Enables a network administrator using a DHCP server supporting DHCP Option 82 to IP addressing policies based on the network area from which a client DHCP request originates. Refer to " DHCP Option 82 " on page 13 .

RADIUS Authentication for Web Browser Access

The following switch models now support use of RADIUS as a primary password authentication method for the Web browser interface:

- Series 2800 (Release I.08.60 and greater)
- Series 2600 and 2600-PWR (Release H.08.58 and greater)
- Series 5300xl (Release E.09.*xx* and greater)

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

Configures RADIUS as the primary password authentication method for the switch's console, Telnet, SSH, and/or the Web browser interface.

`[< local | none >]`

Provides options for secondary authentication (Default: none)

For more on using RADIUS authentication, refer to the chapter titled “RADIUS Authentication and Accounting” in the *Access Security Guide* for your switch.

CLI Local Terminal Mode

This new command enables temporary and non-disruptive changes to the terminal mode without forcing a change in the switch's terminal mode configuration.

Prior to release I.08.60. You could invoke a terminal mode change only with the **console terminal** command, then use **write memory** and reboot the switch. With this older method, you can return the switch to its previous mode only by repeating the whole process. Also, after using **console terminal** to reconfigure the terminal mode, it applies to all console terminal sessions.

Beginning with release I.08.60. This release adds the **console local-terminal** command, which dynamically changes only the console session from which it is executed, does not require **write memory** and a reboot, and does not persist across a reboot.

Syntax: `console local-terminal < vt100 | none | ansi >`

*Dynamically converts the terminal mode of a console session to the selected mode. Executing **console local-terminal** affects only the console session from which it is executed. Rebooting the switch returns the terminal mode for the affected console session to the configured terminal mode. This command does not change the configured console terminal mode configuration. (To change the configured terminal mode, use the **console terminal < vt100 | none | ansi >** command, which requires execution of **write memory**, followed by a switch reboot, to take effect.)*

vt100

*When invoked in a console session, changes the terminal mode to VT-100 for that console session. Use this option when the configured terminal mode is either **none** (scripting mode) or **ansi**, and you want to temporarily use the VT-100 mode. (VT-100 is the default terminal mode configuration setting.)*

none

*When invoked in a console session, changes the terminal mode to “raw” (scripting mode) for that console session. (Scripting mode eliminates unwanted control characters that may appear in some scripting languages.) Use this option when the configured terminal mode is either **vt100** or **ansi**, and you want to temporarily use the scripting mode.*

ansi

*When invoked in a console session, changes the terminal mode to ANSI for that console session. Use this option when the configured terminal mode is either **vt100** (scripting mode) or **none**, and you want to temporarily use the ANSI mode.*

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.

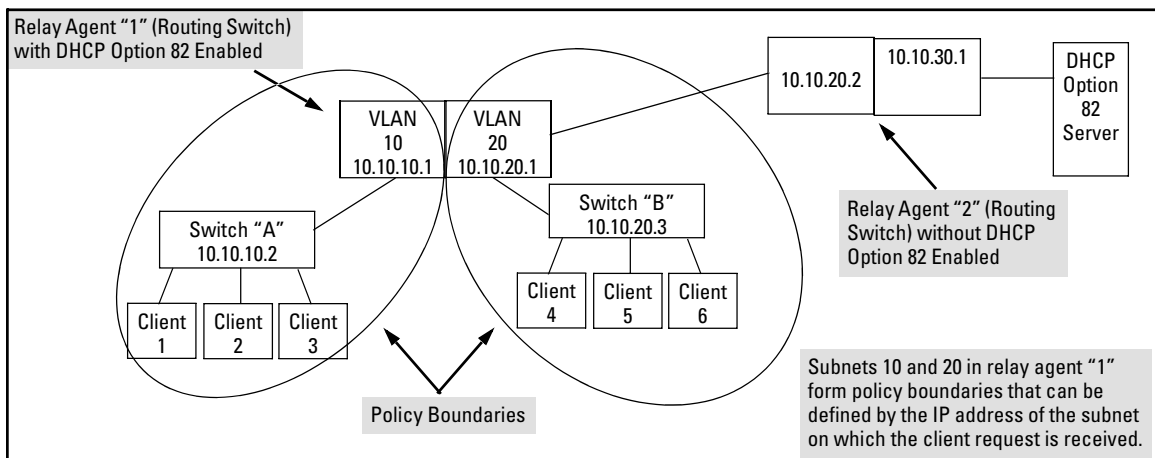


Figure 1. Example of a DHCP Option 82 Application

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 17.)

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 17.)

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.

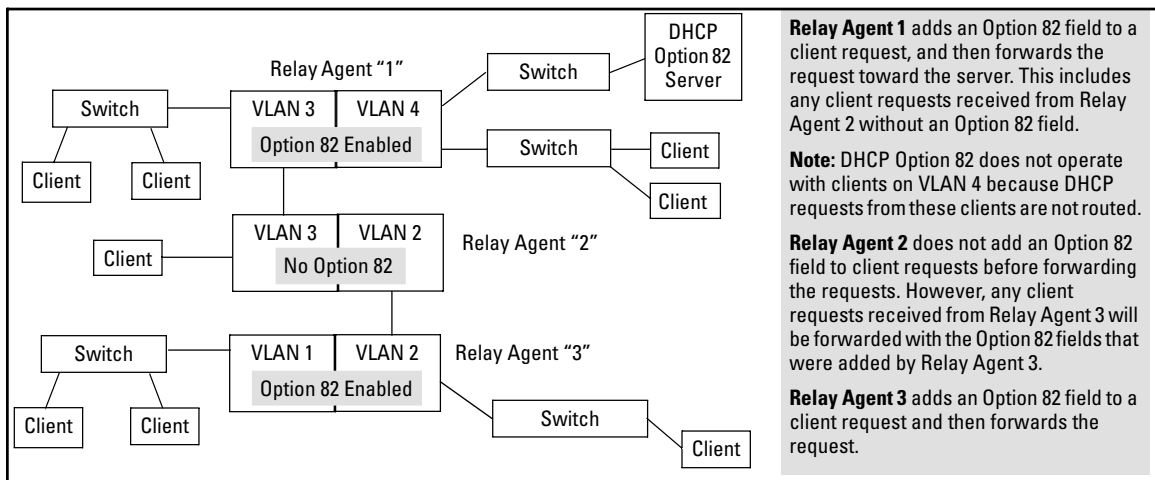


Figure 2. Example of DHCP Option 82 Operation in a Network with a Non-Compliant Relay Agent

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.

```

HPswitch(config)# show system-information
Status and Counters - General System Information

System Name       : HPswitch
System Contact    :
System Location   :

MAC Age Time (sec) : 300

Time Zone         : 0
Daylight Time Rule : None

Firmware revision : I.08.60   Base MAC Addr  : 00110a-a50c20
ROM Version       : I.08.05   Serial Number  : SG426NB048

Up Time          : 32 mins   Memory - Total : 33,043,456
CPU Util (%)     : 4        Memory - Free   : 25,335,136

IP Mgmt - Pkts Rx : 0       Packet - Total  : 1998
          Pkts Tx : 0       Buffers - Free  : 1748
                                   - Lowest   : 1741
                                   - Missed   : 0
  
```

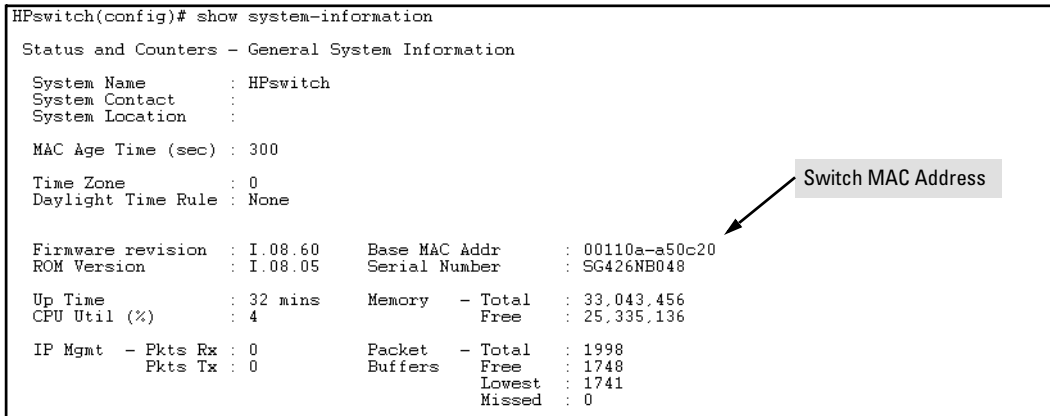


Figure 3. Using the CLI To View the Switch MAC Address

- **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 figure 4, below.)

```
HPswitch(config)# walkmib ifname
ifName.1 = A1
ifName.2 = A2
ifName.3 = A3
ifName.4 = A4
ifName.27 = B1
ifName.28 = B2
ifName.29 = B3
ifName.30 = B4
ifName.31 = B5
ifName.32 = B6
ifName.33 = B7
ifName.34 = B8
ifName.35 = B9
ifName.36 = B10
ifName.37 = B11
ifName.38 = B12
ifName.39 = B13
ifName.40 = B14
ifName.41 = B15
ifName.42 = B16
ifName.43 = B17
ifName.44 = B18
ifName.45 = B19
-- MORE --, next page: Space, next line: Enter, quit: Control-C
```

In this example, the 5304xl has a 4-port module installed in slot “A” and a 24-port module installed in slot “B”. Thus, the first port numbers in the listing are the Index numbers reserved for slot “A”. The first Index port number for slot “B” is “27”, and the Index port number for port B11 (and therefore the Circuit ID number) is “37”.

The Index (and Circuit ID) number for port B11 on a 5304xl routing switch.

Figure 4. 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 1. 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	<p>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.</p> <p>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.</p>
Keep	Append an Option 82 Field	<p>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:</p> <ul style="list-style-type: none"> • 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. <p>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.)</p>
Replace	Append an Option 82 Field	<p>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:</p> <ul style="list-style-type: none"> • 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	<p>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.</p>

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 5. 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 6. 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 7. 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 19.](#)) 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 2, below, illustrates relay agent management of DHCP server responses with optional validation enabled and disabled.

Table 2. 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 drop¹	Drop the server response packet.	Forward server response packet to a downstream device.
	keep²	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 <i>Remote ID</i> and <i>Circuit ID</i> 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 drop¹	Drop the server response packet.	Drop the server response packet.
	keep²	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 <i>Remote ID</i> that did not originate with the relay agent.	append	Drop the server response packet.	Forward server response packet to a downstream device.
	replace or drop¹	Drop the server response packet.	Drop the server response packet.
	keep²	Forward server response packet to a downstream device.	Forward server response packet to a downstream device.
All other server response packets ³	append, keep², replace, or drop¹	Forward server response packet to a downstream device.	Forward server response packet to a downstream device.

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

²A 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).

³A 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 17.*

[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 21.*

[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 17.)

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 17.*

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 I.08.55 Enhancements

Enhancement	Overview
Supports 256 VLANs	Previously, the maximum number of VLANs was 60.
DiffServ Codepoint (DSCP) Marking - L3/L4	Provide support for the following DSCP modes: RFC2474 DiffServ Precedence, RFC2597 DiffServ Assured Forwarding (AF), and RFC2598 DiffServ Expedited Forwarding (EF). (Refer to: Chapter 6, “Quality of Service (QoS): Managing Bandwidth More Effectively on the Series 2600/2600-PWR and Series 2800 Switches” in the <i>Advanced Traffic Management Guide</i> —part number 5990-8853, October 2004— on the ProCurve Networking Web site.*)
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: Chapter 5, “Spanning-Tree Operation” in the <i>Advanced Traffic Management Guide</i> —part number 5990-8853, October 2004— on the ProCurve Networking Web site.*)
Jumbo Packets	On a VLAN configured for jumbo traffic, all ports belonging to that VLAN and operating at 1 Gbps or 10 Gbps allow inbound jumbo packets of up to 9220 bytes (includes four bytes for a VLAN tag). (Refer to: Chapter 10, “Port Status and Basic Configuration” in the <i>Management and Configuration Guide</i> —part number 5990-6023, October 2004— on the ProCurve Networking Web site.*)
Web Authentication	Web authentication adds a new security option that uses a web page login to authenticate users via a RADIUS server for access to the network. (Refer to: Chapter 3, “Web and MAC Authentication for the Series 2600/2600-PWR and 2800 Switches” in the <i>Access Security Guide</i> —part number 5990-6024, October 2004— 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: Chapter 3, “Web and MAC Authentication for the Series 2600/2600-PWR and 2800 Switches” in the <i>Access Security Guide</i> —part number 5990-6024, October 2004— on the ProCurve Networking Web site.*)

Enhancement	Overview (Continued)
MAC Lockdown/Lockout	<ul style="list-style-type: none"> • MAC Lockdown enables the permanent assignment of a MAC address and 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). (Refer to: Chapter 9, “Configuring and Monitoring Port Security” in the <i>Access Security Guide</i>—part number 5990-6024, October 2004—on the ProCurve Networking 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 <i>Management and Configuration Guide</i> —part number 5990-6023, October 2004— on the ProCurve Networking Web site.)*
* To download switch documentation for software release I.08.xx, refer to “To Download Product Documentation:” on page 1.	
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: Chapter 2, “Configuring Username and Password Security” in the <i>Access Security Guide</i> —part number 5990-6024, October 2004— on the ProCurve Networking Web site.)*
Auto-MDI-X manual mode	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. This also allows the manual configuration of port speed. (Refer to: Chapter 10, “Port Status and Basic Configuration” in the <i>Management and Configuration Guide</i> —part number 5990-6023, October 2004— on the ProCurve Networking Web site.)*
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.)
* To download switch documentation for software release I.08.xx, refer to “To Download Product Documentation:” on page 1.	

Release I.07.68 Enhancements

Boot ROM Update

Release I.07.68 is an update to the Boot ROM installer, previously documented in the I.07.64 and I.07.58 software releases. This software installs Boot ROM I.08.05, which aids in preventing some categories of system hang conditions.

Release I.07.64 Enhancements

Boot ROM update

This release (and all subsequent releases) include a Boot ROM upgrade facility. If, during boot, it detects a version 7.xx or a version 8.02 Boot ROM image, it will upgrade the Boot ROM to version I.08.04. This Boot ROM reformats the Flash area to accept code images larger than 3 MB, and is required to be done prior to the next major release of software (expected winter, 2004).

No damage is done to the code stored in the secondary flash area. See the process description under Release I.07.58 below for details.

Release I.07.59 - I.07.63 Enhancements

Software fixes only; no new enhancements.

Includes the Boot ROM upgrade initially introduced in I.07.58 below.

Release I.07.58 Enhancements

Boot ROM update

This release (and all subsequent releases) include a Boot ROM upgrade facility. If, during boot, it detects a version 7.*xx* Boot ROM image, it will upgrade the Boot ROM to version I.08.02. This Boot ROM reformats the Flash area to accept code images larger than 3 MB, and is required to be done prior to the next major release of software (expected winter, 2004).

No damage is done to the code stored in the secondary flash area.

Once the user initially boots with this release (or later I.07.*xx* releases) the console screen will report that the Boot ROM is being updated:

```
Decompressing...done.

Initializing...

WARNING : This release includes a BootROM upgrade.
          Interrupting this process will cause
          The switch to become unusable.

BootROM upgrade in progress...completed.
```

This process takes approximately 3 minutes to complete, including the normal time to boot up your switch. Note: previous versions of these release notes placed this Boot ROM upgrade facility in Release I.07.56.

Release I.07.53 - I.07.57 Enhancements

Software fixes only; no new enhancements.

Release I.07.52 Enhancements (Beta Only)

QOS Pass-Through Mode

Release I.07.52 introduced a new command to enhance the performance of line-rate traffic transfers through the 2800 Series switches. This feature should only be used in environments where Quality of Service (QoS) is not of major importance, but where lossless data transfers are key. This command essentially disables any discrimination of QoS queues for traffic, consolidating packet buffer memory to provide line-rate flows with no loss of data.

For more information, refer to the chapter titled “Port Status and Basic Configuration” in the *Management and Configuration Guide* for your switch (October 2004 version, or later). To download the latest version of switch documentation, refer to [“Downloading Switch Documentation and Software from the Web”](#) on page 1.

Release I.07.51 Enhancements

Software fixes only; no new enhancements.

Release I.07.50 Enhancements

Port Trunking

New Maximum for Number of Ports in a Trunk

Trunk groups can now be configured with up to 8 ports per trunk. (Formerly, the switches allowed only 4 ports per trunk.) Also, you can now configure up to 24 trunk groups per switch. (Formerly, the switches allowed only 6 port trunk groups).

Restriction on Grouping of Ports in a Trunk with IP Routing Enabled

Trunk groups can generally be specified as any grouping of ports on the switch. However, if IP routing is enabled on the switch, all of the ports in a given trunk group must be in the same range of ports. These ranges are as follows:

- 1-12
- 13-24
- 25-36 (applies only to the Switch 2848)
- 37-48 (applies only to the Switch 2848)

For more information, refer to the chapter titled “Port Trunking” in the *Management and Configuration Guide* for your switch (October 2004 version, or later). To download the latest version of switch documentation, refer to [“Downloading Switch Documentation and Software from the Web”](#) on page 1.

Port Monitoring

In software releases prior to release I.07.50, port monitoring sent only inbound (ingress) traffic to the monitor. Beginning with release I.07.50, the Series 2800 switches will now also send outbound (Egress) traffic to the mirror port when port monitoring is enabled.

Release I.07.32 Enhancements

Software fixes only; no new enhancements.

Software Fixes in Release I.07.xx - I.08.xx

Release I.07.31 was the first software release for the ProCurve 2800 Series.

Release I.08.62

Problems Resolved in Release I.08.62

- **Crash (PR_1000207542)** — The switch may crash with a bus error or a task hang.
- **Crash (PR_1000215009)** — Software exception in ISR at `intr.c:595 -> FATAL SCHAN ERROR`.
- **Enhancement (PR_1000213492)** — QoS-Passthrough can be enabled at run time vs. requiring a write memory and reboot. (Enhancement documentation to be provided at a later time.)
- **Flow Control (PR_1000217576)** — Flow control/jumbo frame error messages not generated in event log.
- **Port Security (PR_1000203984)** — CLI port-security command - mac-address command will save more addresses than is configured.

Release I.08.61

Problems Resolved in Release I.08.61

- **CLI (PR_1000214598)** — The switch does not accept the CLI command "spanning-tree 1 mode fast".
- **Config (PR_1000216051)** — Reloading a previously saved startup-configuration with command "stack join (mac address)" to a member switch of the IP stack breaks the membership of that same stack. Commander hangs with member "mismatched".
- **Web (PR_80857)** — A problem with IE4 and WebAgent. Recompiled the Web Agent with a new Java Development Kit (1.2 - was 1.1)

Release I.08.60

Problems Resolved in Release I.08.60

- **ACL (PR_1000207620)** — The switch sometimes incorrectly permits TCP and UDP traffic in spite of an ACL configuration.
- **CLI (PR_1000202435)** — When IGMP fast-leave is configured via the CLI, the configuration is not displayed with the "show configuration" command.
- **Config (PR_1000087886)** — The CLI will display error message "Value 1000-full is not applicable to port <port num>", when trying to download a startup-configuration with a Mini-GBIC module configured at 1000 full duplex.
- **Crash (PR_1000205768)** — "null" System Name in the Web user interface may crash with: "Software exception at lldpSysNameTlv.c:251 – in 'mlldpCtrl', >task ID = 0x12dc88 -> ASSERT: failed".
- **Crash (PR_1000200341)** — In some cases a protocol or feature may not function correctly.
- **Crash (PR_1000208530)** — Unpredictable results
- **Crash (PR_1000201614)** — When the switch is set with a 16 character manager password within the setup menu, a 'Bus error' crash may occur.
- **DHCP Enhancement (PR_1000207639)** — DHCP Option 82 implementation (DHCP Tracker).
- **DHCP Relay (PR_1000207419)** — The DHCP Relay agent was disabled by default in earlier Version 8 releases. With this fix, the DHCP Relay agent is enabled by default, as it was in I.07 releases.
- **IP Helper/DHCP Relay (PR_1000197046)** — The switch may not handle "DHCP Inform" relay messages properly from the client, resulting in a failed transaction.
- **Management enhancement (No PR)** — Non-Persistent console terminal mode.
- **Other PR_1000209839)** — Memory corruption of dmaStats do to off array boundary error.
- **Other PR_1000200341** — Added an exception handler to prevent a case where the system may hang.
- **Open VLAN (PR_1000210932)** — Open VLAN mode (Unauthorized VLAN) does not work with any Port-Security Learn-Mode.
- **RMON (PR_1000196477)** — When RMON thresholds in the switch are exceeded no trap is generated.

- **SNMP (PR_1000196170)** — Traps are not buffered before the IP stack is initialized, causing the possibility of missing some traps generated during startup.
- **SNMP (PR_1000212170)** — The Switch transmits Warm and Cold Start traps with an agent address of 0.0.0.0.
- **Testmode (PR_1000212159)** — Added the testmode command 'memWatch'
- **Web Enhancement (NO PR)** — RADIUS for the Web browser interface.
- **Web UI/Port Security (PR_1000195894)** — The Web user interface does not allow the user to select multiple ports when configuring port-security.
- **Web UI (PR_1000191635)** — The Port column may not be sorted correctly in all Web user interface screens.
- **Web UI (PR_93721)** — Scroll bar does not work in Web Status screen.
- **Web UI (PR_1000210110)** — Slow Web UI performance.

Release I.08.58

Problems Resolved in Release I.08.58

- **802.1s (PR_1000207608)** — After the root bridge is agreed, the non-root switch continues to send out BPDUs claiming to be Root, resulting in possible instability in the STP topology.

Release I.08.57

Problems Resolved in Release I.08.57

- **Port Hang (PR_1000212920)** — Unpredictable switching and LED behavior where one or more ports may cease to forward traffic.
- **SNMP (PR_1000190654)** — Some of the fault finder events in the SNMP traps list a 0.0.0.0 IP address in the URL.

Release I.08.56

Problems Resolved in Release I.08.56

- **Port Hang (PR_1000207174)** — Reduces or eliminates the occurrence of "port hang" issues, where one or more ports may cease to forward traffic and the LEDs display status may be incorrect.

Release I.08.55

Problems Resolved in Release I.08.55

- **CLI PR_82258** — **sh ip igmp** command shows blank lines inter-mixed within the displayed table.
- **CLI PR_1*3169** — 2800: "port-security learn-mode configured " is shown as "static" in CLI.
- **CLI PR_1*11958** — Add CLI command to configure the outbound queue (2 or 4 queue).
- **CLI PR_1*18700** — **Show ip route** "IP Route Entries" not centered in output.
- **Config PR_92346** — Unable to delete empty VLAN.
- **Crash PR_91463** — Displays a crash message similar to the following:
Software exception at ip_util.c:413 -- in 'ifInfo', task ID = 0x143fb60
- **Crash PR_93791** — Displays a crash message similar to the following:
Software exception at bcmHwFeatures.c:108 -- in 'mAdMgrCtrl'
- **Crash PR_1*1537** — Memory leak on 2848, fatal exception in malloc_else_fatal().
- **Crash PR_1*5466** — Displays a crash message similar to the following:
Software exception in ISR at bcm56xxDmaPoll.c:623 (top-of-tree)
- **Crash PR_1*20805** — Displays a crash message similar to the following:
Software exception @ route.c:331 (attempting to free an already freed rtenry)
- **Crash PR_1*21853** — Displays a crash message similar to the following:
Software exception @ radix.c:922, route does not exist in the tree
- **Help PR_98206** — Help file is not consistent with the actual usage.
- **Help PR_1*21395** — Help text incorrect for some ip icmp commands.
- **Hot Swap PR_1*18578** — Dual personality ports on 2800 and 2600 have hotswap out problem.
- **LACP PR_1*6404** — Dynamic LACP: Standby mode problem.
- **MCAST PR_1*6552** — Multicast pkts flooded on a VLAN w/ igmp enabled; hw & sw out of sync.
- **Menu PR_94905** — 2848 Configuration screen shows only 32 VLANs of configured 48.

- **Ping PR_1*19945** — Unable to ping through default gateway.
- **Routing PR_1*5961** — Layer 3 connectivity lost when address is moved across ASIC port group.
- **Routing PR_1*20234** — 2800 / I.07.53 / I.07.52: DD IGMP Squelches EIGRP when triggered by SSDP Packet.
- **SNMP PR_88716** — SNMP walk times out with large configuration.
- **SNMP PR_1*3361** — 'snmpv3' configtest failure.
- **Syslog PR_97016** — syslog word-complete options are not consistent between 6108 and 2800.
- **VLAN PR_90884** — VLAN PORT_UNTAGGEDMAP config not being set correctly.
- **VLAN PR_92413** — Broadcasting is forwarded outside the VLAN.
- **Web PR_1*1216** — Web UI, log error.
- **Web PR_1*12103** — Garbage in the Web UI Status | Overview screen.
- **Web PR_1*21294** — Stack Management Screen is blank.
- **Web PR_1*21867** — Web UI VLAN Configuration is broken.

Release I.07.68

Problems Resolved in Release I.07.68

- **Other PR_1000200341**— Added an exception handler to prevent a case where the system may hang.

Release I.07.67

Problems Resolved in Release I.07.67

- **Port Hang (PR_1000212920)** — Contains and automatically installs (after reboot) Boot ROM version I.08.05. This Boot ROM fixes the 'port hang' issue which can result in unpredictable switching and LED behavior.

Release I.07.66

Problems Resolved in Release I.07.66

- **Dead Port PR_1*207174** — Reduces or eliminates the occurrences of the "dead-port" issue. "Dead-port" implies the following symptoms:
 - Some 2800 ports may not forward packets, while other ports continue to forward packets.
 - Link LED stays on when cable removed
 - Link LED stays off when cable attached.

Release I.07.65 (Not Released)

Problems Resolved in Release I.07.65

- **Other** — Added diagnostic code for isolating the "dead-port" issues at hot-site.

Release I.07.64

Problems Resolved in Release I.07.64

- **Boot ROM PR_1*202277** — Contains and installs the I.08.04 Boot ROM which fixes the 'port hang' problem when back-revving to software versions older than I.07.58 when running the I.08.03 Boot ROM. Also, improves on the 'reload hang' fix that is in I.08.03 by specifically addressing the 'boot hang' that can occur after the I.08.03 Boot ROM's patcher (the 'patcher' is the part of the software that installs the new Boot ROM) does a reboot.

Release I.07.63 (Beta Only)

- **Boot ROM PR_1*85713** — Introduced the I.08.03 Boot ROM patcher to address a 'boot hang' issue. This Boot ROM was later updated to I.08.04 (see Release I.07.64).

Release I.07.62 (Beta Only)

- **QOS PR_1*194538** — QOS-Pass -Through-Mode (introduced in I.07.52) does not work in software versions I.07.54-I.07.61

Release I.07.61

Problems Resolved in Release I.07.61

- **DHCP Relay PR_1*188635** — DHCP Relay sometimes preserves the incoming MAC SA in relayed packets.
- **HANG/WEB PR_1*190109** — Fix for cases where the Web interface would stop responding when the user enters the Configuration Screen. Once triggered, no access to the Web agent is possible from any client.

Release I.07.60

Problems Resolved in Release I.07.60

- **Auto-TFTP/Rebooting PR_1*20802** — Auto-TFTP causes constant rebooting, with no resulting crash files.
- **Auto-TFTP PR_1*187649** — Auto-TFTP will not allow a forced download of software after Auto-TFTP is Disabled.
- **Hang PR_1*190119** — Additional case of system hang found and addressed.

Release I.07.59 (Beta Only)

Release I.07.58 (Beta Only)

Problems Resolved in Release I.07.58

- **IGMP PR_1*06552 and 1*20234** — The switch floods multicast packets on a VLAN when IGMP is enabled, due to h/w & s/w MAC tables being out of sync.
- **TELNET PR_1*19573** — Switch reboots when telnet is disabled and port 1506 accessed. The switch produces no crash-log.
- **Web PR_89899** — In the Web UI, port statistic counters are overwriting one another.
- **IGMP/EIGRP PR_1*20234** — With IGMP enabled the switch drops EIGRP packets (when triggered by receiving an SSDP packet).
- **VLAN PR_95593** — The switch will not allow the user to delete a VLAN that contained a mini-GBIC port that was removed.

- **CLI/Config PR_1*01628** — In the CLI, switch reports “Inconsistent value” error when adding ports to a VLAN.
- **Hot-swap/Config PR_1*89150** — Switch configuration is not properly updated on transceiver swap events

Release I.07.57 (Never Released)

Release I.07.56

Problems Resolved in Release I.07.56

- **Hang PR_1*87409, PR_1*6985** — Symptoms vary, and can include any of the following:
 - Switch does not respond to pings, WEB access, Telnet, or Console access.
 - Pre-existing links prior to the “hang” still appear to transmit and receive data normally.
 - LED behavior on ports that establish link before the “hang” is erratic. For example, the LED remains lit even after dropping physical link.
 - New links attempted after the “hang” do not transmit or receive traffic. Also that port's LED on the Switch 2800 remains dark, indicating no link while the neighbor device's LED may light up indicating that link is established.
 - Front panel LED Mode, Reset, and Clear buttons, may not function properly.
 - CPU-dependent features such as STP may not function properly.

A power cycle of the switch has been the only way to relieve these “hang” symptoms. Since the switch agent does not respond and the front panel buttons may not respond, it may be necessary to unplug and re-plug the power cable in order to reset the switch.

Release I.07.55 (Beta Only)

Problems Resolved in Release I.07.55

- **Crash PR_1*20824** — Displays a crash message similar to the following:

```
SubSystem 0 went down: 01/02/90 22:33:36 NMI occurred: IP=0x003164b0
MSR:0x0000b032 LR:0x003164d4 Task='tDPC' Task ID=0x1ad2440 cr:
0x28000080 sp:0x01ad2380 xer:0x00000000.
```
- **Flow Control PR_98957** — The switch honors PAUSE (flow control) frames that it receives, but it does not generate them.
- **Show Mac PR_82086** — The CLI command **show mac** < mac-address > does not work.

- **OpenSSL/crash PR_1*12823** — OpenSSL bus error vulnerability.

Release I.07.53 (Beta Only)

Problems Resolved in Release I.07.53

- **Crash PR_1*3390** — Memory leak causing crash `sw_malloc.c:141` in `snmpevt` task.
- **Crash PR_1*13156** — Master crash in memory system - `memPartFree`. The specific crash symptoms can vary widely.
- **RMON PR_1*11690** — The switch does not send RMON trap PDUs.

Release I.07.52 (Beta Only)

Problems Resolved in Release I.07.52

- **GVRP PR_1*5082** — Vague error message (`commit failed`) when trying to add more than the maximum number of allowed VLANs.
- **Performance PR_1*11958** — Enhancement: Added the **qos-passthrough-mode** configuration option to the CLI to configure the number of outbound queues to use. Refer to “[QOS Pass-Through Mode](#)” on page 28.
- **Trunking PR_1*5962** — Unable to form LACP dynamic trunk across ASIC port groups without routing enabled.
- **sysUptime PR_1*4025** — `sysUptime` wraps in approximately 49 days.
- **Web PR_1*4111** — The Stack Management view has a scroll problem.
- **Web PR_1*3580** — The web interface allows broadcast and multicast destination addresses.
- **Web PR_1*7144** — VLAN Configuration Help link is not available.

Release I.07.50

Problems Resolved in Release I.07.50

- **CLI PR_97671** — If the number of `max-vlans` is greater than 15 and the user tries to add new vlan the Switch reports `Commit failed`. In the Web user interface the Switch reports An error was encountered while attempting to add the VLAN entry.

The message is changed to:

Maximum number of VLANs (max-vlans) has already been reached.

■ **CLI PR_1*3517** — Counters. Various, related issues:

- ifInDiscards (RX drops in the menu interface) includes outbound drops; fixed to display only true inbound drops.
- ifOutDiscards does not report out-bound drops; fixed to display true outbound drops previously shown on ifInDiscards.
- dot1dtpPortInDiscards includes outbound drops; fixed to display only inbound discards.
- ipInDiscards includes outbound drops; fixed to show only inbound IP based discards.

■ **Crash PR_95525** — Various crashes, including:

```
Bus error: HW Addr=0xe1f08796 IP=0x003a51b4 Task='mInstCtrl' Task ID=0x1767af8 fp: 0x00000006 sp:0x01767988 lr:0x003979a4
```

■ **Crash PR_1*2979** — Software exception at rstp_port_role_sm.c:44 -- in mRstpCtrl.

■ **GVRP PR_1*3124** — Uncertain error message when trying to add more than max VLANs.

■ **Port Monitoring PR_1*3540 Enhancement** — Add Egress (output) port monitoring.

■ **RSTP PR_1*1612** — Under some circumstances a port may take approximately 30 seconds to go into Forwarding state.

■ **Services PR_1*3867** — ICMP Redirects never age. Causes any incoming or outgoing agent communications such as ping, TELNET, Web, SNMP, etc. to fail with a message similar to the following:

```
HW Addr=0x000-0000-0 IP=0x0-02a22d8 Task='tNetTask' Task ID=0xe2e740.
```

■ **Trunks PR_1*3530** — Enhancement: Increase the limit on trunks and ports per trunk to:

- Up to 24 trunks, total; and
- Up to 8 ports per trunk

Refer to [“New Maximum for Number of Ports in a Trunk” on page 28](#).

■ **Web/IP Stack Management PR_92826** — With an eight switch IP Stack Management stack, management of the switches with the Web interface can cause the commander switch to crash or hang. If the user selects options too quickly or moves from one option to another the Web user interface may freeze and become unresponsive. The Commander Switch may also crash with a Bus Error. Also, TELNET and Console interfaces may become unresponsive.

■ **Web/IP Stack Management PR_97323** — In the Web user interface the images displayed for the stack members are not correct.

- **Web PR_98500** — Clicking on tabs in a certain order causes the browser window to close (terminate).
- **Web/SSL PR_98918** — When creating an SSL certificate the Organization name and unit are switched in the web user interface display. **Emphasis: This is only a display issue.**
- **Web PR_81848** — **The [Clear changes] button does not work for the Default Gateway or VLAN selections**
- **Web PR_82199** — VLAN port modification shows misleading mode. In the Configuration - VLANs - Modify page, select a port, then set the “mode” modify pull-down menu to “tagged”. Select another port. The “mode” pulldown field remains set to “tagged”, which is misleading and incorrect, in general.
- **Web PR_97407** — Port security error message is unclear with mac lockdown. The user interface may report “**Unable to add new MAC Address. MAC entry is either a multicast, broadcast or NULL address.**” when, in fact, the MAC address the user is specifying is locked down or locked out.
- **Web PR_1*452** — Resetting the Switch leads to the URL **aol.co.uk**.
- **Web PR_90858** — VLAN Name text field won't clear after 12 characters are entered.
- **Web PR_1*1702** — Sometimes clicking on the **[Apply]** button on the Configuration/Monitor Port screen results in the message **Not enough params specified**.
- **Web PR_92078** — After making changes under the Device Features tab, the page never fully loads.
- **Web PR_82039** — If the user selects GVRP mode, selects a port, and then selects nothing as an option for the port mode, all ports below the selected port disappear. This does not affect the switch configuration.

Release I.07.32

Problems Resolved in Release I.07.32

- **Command Line Interpreter PR_95284** — A too long MAC addresses in a port-security CLI command results in:

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

Here is an example command that would have crashed Version I_07.31:

```
port-security 1 learn-mode static address-limit 1 mac-address 0800000000010000000
```

- **SSH PR_96648** — CERT Advisory CA-2003-24: OpenSSH vulnerability. Fix implemented. For details, see CERT Advisory CA-2003-24 and associated vulnerability note “VU#333628” at <http://www.cert.org/advisories/CA-2003-24.html>.
- **System Log PR_95689** — Excessive Time Sync entries when using a Timep or SNTP server. The system software needed to be adjusted to properly keep synchronized with a configured SNTP server. In earlier versions of software, this resulted in an excessive number of Time Sync entries in the event log. This only applies to the 2800 Series switches running I_07.31 software.

Known Software Issues and Limitations

Issues

None at this time.

Limitations

Displaying the Fast-Leave Setting on a Port

Use the **walkmib** command, below, to display this setting for all switch ports or the ports on a specified VLAN.

Syntax:

```
walkmib hpSwitchIcmpPortFastLeaveState<.vlan number>
```

```
HPswitch(vlan-1)# walkmib hpswitchicmpportfastleavestate.1
hpSwitchIcmpPortFastLeaveState.1.1 = 1
hpSwitchIcmpPortFastLeaveState.1.2 = 2
hpSwitchIcmpPortFastLeaveState.1.3 = 1
hpSwitchIcmpPortFastLeaveState.1.4 = 2
hpSwitchIcmpPortFastLeaveState.1.5 = 1
hpSwitchIcmpPortFastLeaveState.1.6 = 2
```

The **2** at the end of a port listing shows that Fast-Leave is **disabled** on the corresponding port.

The **1** at the end of a port listing shows that Fast-Leave is **enabled** on the corresponding port.

Internal VLAN Number for the Default VLAN
Note: Internal VLAN numbers reflect the sequence in which VLANs are created, and are not related to the unique VID assigned to each VLAN.

Sequential Port Numbers (not all ports shown here)

— *This page is intentionally unused.* —



© 2001, 2005 Hewlett-Packard Development
Company, LP. The information contained
herein is subject to change without notice.

April 2005
Part Number
5990-6049