

Release Notes: Version M.10.72 Software for the HP ProCurve Series 3400cl Switches

"M" software versions are supported on these switches:

ProCurve Switch	M.08.51 through M.08.95	M.08.99.x and newer	M.08.96, M.08.97, M.10.01 and newer
ProCurve Switch 3400cl-24G (J4905A)	✓		✓
ProCurve Switch 3400cl-48G (J4906A)	✓		✓
ProCurve Switch 6400cl-6XG 10-GbE CX4(J8433A)	✓	✓	
ProCurve Switch 6410cl-6XG 10-GbE X2(J8474A)	✓	✓	

Release M.10.41 supports the ProCurve Switch 3400cl-24G (J4905A), and 3400cl-48G (J4906A). These release notes include information on the following:

- Downloading switch software and documentation from the Web ([page 1](#))
- Clarification of operating details for certain software features ([page 20](#))
- A listing of software enhancements in recent releases ([page 25](#))
- A listing of software fixes included in releases M.08.51 through M.10.72 ([page 145](#))

IMPORTANT:

3400cl switches MUST be running ROM version I.08.12 prior to loading M.10.20 or newer software. If your switch is using a software version earlier than M.10.10, you need to install and boot the M.10.10 software (included in the M.10.41 software package) to load the I.08.12 ROM version, before installing M.10.20 or newer.

Security Note:

Downloading and booting software release M.08.89 or greater for the first time automatically enables SNMP access to the hpSwitchAuth MIB objects. If this is not desirable for your network, ProCurve recommends that you disable it after downloading and rebooting with the latest switch software. For more information, refer to "Enforcing Switch Security" on page 10 and "Using SNMP To View and Configure Switch Authentication Features" on page 35.

Configuration Compatibility Caution:

Configuration files created or saved using version M.10.65 or higher are NOT backward-compatible with previous software versions. The user is advised to save a copy of the pre-M.10.65 startup-config file **BEFORE UPDATING** to M.10.68 or greater, in case there is ever a need to revert back to an earlier version of software.

© Copyright 2004 - 2009 Hewlett-Packard Development Company, LP. The information contained herein is subject to change without notice.

Publication Number

5991-4764
May, 2009

Applicable Product

ProCurve Switch 3400cl-24G	(J4905A)
ProCurve Switch 3400cl-48G	(J4906A)

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	1
Software Updates	1
Download Switch Documentation and Software from the Web	1
View or Download the Software Manual Set	1
Downloading Software to the Switch	1
Downloading Software to the Switch	2
TFTP Download from a Server	3
Xmodem Download From a PC or Unix Workstation	3
Saving Configurations While Using the CLI	5
Install Recommendations for I.08.12 Boot ROM Update	6
ProCurve Switch, Routing Switch, and Router Software Keys	7
Minimum Software Versions for Series 3400cl Switch Features	9
OS/Web/Java Compatibility Table	9
Enforcing Switch Security	10
Switch Management Access Security	10
Default Settings Affecting Security	10
Local Manager Password	11
Inbound Telnet Access and Web Browser Access	11
Secure File Transfers	11
SNMP Access (Simple Network Management Protocol)	12
Physical Access to the Switch	13
Other Provisions for Management Access Security	14
Network Access Security	15
Access Control Lists (ACLs)	15
Web and MAC Authentication	15
Secure Shell (SSH)	16
Secure Socket Layer (SSLv3/TLSv1)	16
Traffic/Security Filters	16
802.1X Access Control	17
Port Security, MAC Lockdown, MAC Lockout, and IP Lockdown	18
Key Management System (KMS)	18

Connection-Rate Filtering Based On Virus-Throttling Technology	19
Identity-Driven Management (IDM)	19
Clarifications and Updates	20
Operating Notes for Jumbo Traffic-Handling	20
Non-Genuine Mini-GBIC Detection and Protection Initiative	20
Publication Updates	20
IGMP Command Update	21
General Switch Traffic Security Guideline	22
The Management VLAN IP Address	23
Interoperating with 802.1s Multiple Spanning-Tree	23
Rate-Limiting	23
Known Issues	24
Release M.10.17	24
Enhancements	25
Release M.08.69 Enhancements	25
Release M.08.70 through M.08.72 Enhancements	25
Release M.08.73 Enhancements	25
Release M.08.74 through M.08.77 Enhancements	25
Release M.08.78 Enhancements	26
Using Fastboot To Reduce Boot Time	26
Release M.08.79 Enhancements	26
CLI Port Rate Display	26
Release M.08.80 through M.08.83 Enhancements	27
Release M.08.84 Enhancements	28
Release M.08.85 through M.08.88 Enhancements	28
Release M.08.89 Enhancements	28
DNS Resolver	28
Using SNMP To View and Configure Switch Authentication Features	35
Releases M.08.90 and M.08.91 Enhancements	38
MSTP Default Path Cost Controls	38

QoS Pass-Through Mode	39
Release M.08.94 Enhancements	42
DHCP Option 82: Using the Management VLAN IP Address for the Remote ID	42
UDP Broadcast Forwarding	44
Releases M.08.95 through M.10.01 Enhancements	45
Release M.08.96 Enhancements	45
Releases M.08.97 through M.10.01 Enhancements	45
Release M.10.02 Enhancements	45
RADIUS-Assigned Access Control Lists (ACLs)	45
SFlow Show Commands	68
Release M.10.04 Enhancements	70
Instrumentation Monitor	70
TCP/UDP Port Closure	75
Spanning Tree Show Commands	77
Release M.10.05 Enhancements	79
Release M.10.06 Enhancements	79
Release M.10.07 Enhancements	80
Release M.10.08 Enhancements	80
Release M.10.09 Enhancements	80
Uni-Directional Link Detection (UDLD)	80
Release M.10.10 Enhancements	88
Spanning Tree Per-Port BPDU Filtering	88
Releases M.10.11 through M.10.12 Enhancements	91
Release M.10.13 Enhancements	91
Releases M.10.14 through M.10.16 Enhancements	91
Release M.10.17 Enhancements	91
Spanning Tree BPDU Protection	91
Example of BPDU Protection Additions to Show Spanning Tree Command	94
Release M.10.21 Enhancements	95
Release M.10.22 Enhancements	95
Release M.10.23 Enhancements	97
Release M.10.24 Enhancements	97
Release M.10.25 Enhancements	97

Release M.10.26 Enhancements	97
Release M.10.27 Enhancements	98
Release M.10.28 Enhancements	100
Release M.10.29 Enhancements	100
Release M.10.30 Enhancements	100
Release M.10.31 Enhancements	100
Release M.10.32 Enhancements	101
Scheduled Reload	101
Release M.10.33 Enhancements	102
How RADIUS-Based Authentication Affects VLAN Operation	102
VLAN Assignment on a ProCurve Port	102
Operating Notes	103
Example of Untagged VLAN Assignment in a RADIUS-Based Authentication Session	104
Enabling the Use of GVRP-Learned Dynamic VLANs in Authentication Sessions	107
Release M.10.34 Enhancements	108
Concurrent TACAS+ and SFTP	108
Release M.10.35 Enhancements	109
Dynamic ARP Protection	109
Release M.10.36 Enhancements	115
Release M.10.37 Enhancements	115
Configuring MSTP Port Connectivity Parameters	116
Release M.10.38 Enhancements	118
Send SNMP v2c Informs	119
Release M.10.39 Enhancements	120
RADIUS Server Unavailable	121
ARP Age Timer Increase	124
Release M.10.40 Enhancements	126
Release M.10.41 Enhancements	126
Release M.10.42 Enhancements	126
Release M.10.43 Enhancements	126
Dynamic IP Lockdown	126
Operating Notes	130
Release M.10.44 through M.10.64 Enhancements	135

Release M.10.65 Enhancements	136
MSTP VLAN Configuration Enhancement	136
Release M.10.66 Enhancements	140
Configure Logging via SNMP	140
Release M.10.67 Enhancements	143
Release M.10.68 Enhancements	143
LACP and Link Traps Global Disable	143
Release M.10.69 Enhancements	144
Release M.10.70 Enhancements	144
Release M.10.71 Enhancements	144
Release M.10.72 Enhancements	144
Software Fixes in Release M.08.51 - M.10.72	145
Release M.08.52	145
Release M.08.53 (Never Released)	145
Release M.08.54	145
Release M.08.55 - Release M.08.60	145
Release M.08.61	145
Release M.08.62	147
Release M.08.63	147
Release M.08.64	148
Release M.08.65	148
Release M.08.66	148
Release M.08.67	148
Release M.08.68	149
Release M.08.69	149
Release M.08.70	150
Release M.08.71	150
Release M.08.72	151
Release M.08.73	151
Release M.08.74	151
Release M.08.75	152

Release M.08.76	152
Release M.08.77	152
Release M.08.78	153
Release M.08.79	153
Release M.08.80	153
Release M.08.81	153
Release M.08.82	153
Release M.08.83	154
Release M.08.84	154
Release M.08.85	154
Release M.08.86	154
Release M.08.87	155
Release M.08.88	155
Release M.08.89	155
Release M.08.90	156
Release M.08.91	156
Release M.08.92	156
Release M.08.93	157
Release M.08.94	157
Release M.08.95	157
Release M.08.96	157
Release M.08.97	158
Release M.10.01	158
Release M.10.02	158
Release M.10.03	158
Release M.10.04	159
Release M.10.05	159
Release M.10.06	159
Release M.10.07	160
Release M.10.08	160
Release M.10.09	161

Release M.10.10	161
Release M.10.11	162
Release M.10.12	162
Release M.10.13	162
Release M.10.14	163
Release M.10.15	163
Release M.10.16	163
Release M.10.17	164
Release M.10.18 - Release M.10.19	164
Release M.10.20	164
Release M.10.21	165
Release M.10.22	165
Release M.10.23	166
Release M.10.24	166
Release M.10.25	166
Release M.10.26	167
Release M.10.27	167
Release M.10.28	168
Release M.10.29	168
Release M.10.30	169
Release M.10.31	169
Release M.10.32	170
Release M.10.33	170
Release M.10.34	171
Release M.10.35	171
Release M.10.36	172
Release M.10.37	172
Release M.10.38	172
Release M.10.39	173
Release M.10.40	173
Release M.10.41	173

Release M.10.42	174
Release M.10.43	174
Release M.10.44	174
Release M.10.45	175
Release M.10.46	175
Release M.10.47	175
Release M.10.48	176
Release M.10.49	176
Release M.10.50 through M.10.64	177
Release M.10.65	177
Release M.10.66	178
Release M.10.67	179
Release M.10.68	180
Release M.10.69	180
Release M.10.70	181
Release M.10.71	183
Release M.10.72	183

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.

Download Switch Documentation and Software from the Web

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

View or Download the Software Manual Set

Go to: www.procurve.com/manuals

You may want to bookmark this Web page for easy access in the future.

You can also register on the My ProCurve portal to receive a set of ProCurve switch manuals on CD-ROM. To register and request a CD, go to www.procurve.com and click on **My ProCurve Sign In**. After registering and entering the portal, click on **My Manuals**.

Downloading Software to the Switch

ProCurve Networking periodically provides switch software updates through the ProCurve Networking Web site (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:
 - Select **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:
 - Select **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 3).
- Use the USB port to download a software file from a USB flash drive.
- Use the download utility in ProCurve Manager Plus.

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.

Downloading Software to the Switch

ProCurve Networking periodically provides switch software updates through the ProCurve Networking Web site (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 3).
- Use the download utility in ProCurve Manager Plus.

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 M_08_8x.swi from a TFTP server with the IP address of 10.28.227.103:

1. Execute the copy command as shown below:

```
ProCurve switch # copy tftp flash 10.28.227.103 M_08_8x.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 shown in [Figure 1](#). When the CLI prompt re-appears, the switch is ready to reboot to activate the downloaded software:

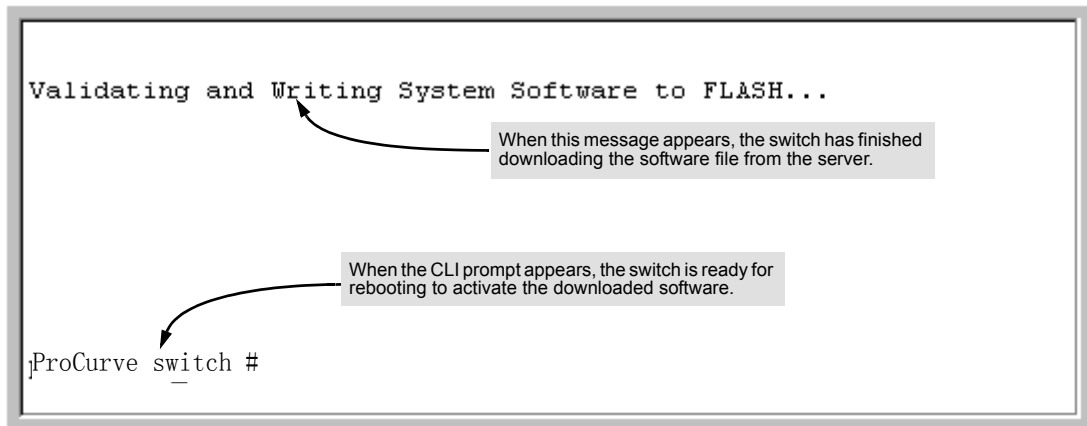


Figure 1. Message Indicating the Switch Is Ready To Activate the Downloaded Software

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 to a PC operating as a terminal. (Refer to the Installation and Getting Started Guide you received with the switch for information on connecting a PC as a terminal and running the switch console interface.)
- The switch software is stored on a disk drive in the PC.

Software Management

Downloading Software to the Switch

- The terminal emulator you are using includes the Xmodem binary transfer feature. (For example, in the HyperTerminal application included with Windows NT, you would use the Send File option in the Transfer dropdown menu.)

Using Xmodem and a terminal emulator, you can download a switch software file to either primary or secondary flash using the CLI.

Syntax: `copy xmodem flash [< primary | secondary >]`

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 115200 bits per second. (The baud rate must be the same in both devices.) For example, to change the baud rate in the switch to 115200, execute this command:

```
ProCurve (config)# console baud-rate 115200
```

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

Changing the console baud-rate requires saving to the Startup Config with the "write memory" command. Alternatively, you can logout of the switch and change your terminal emulator speed and allow the switch to AutoDetect your new higher baud rate (i.e. 115200 bps)

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. For example, using HyperTerminal:
 - a. Click on Transfer, then Send File.
 - b. Type the file path and name in the Filename field.
 - c. In the Protocol field, select Xmodem.
 - d. Click on the Send button.

The download can take several minutes, depending on the baud rate used in the transfer.

4. 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.)
5. 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.

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] ?
```

Install Recommendations for I.08.12 Boot ROM Update

When installing the M.10.17 software to load the I.08.12 ROM version, ProCurve recommends that you use the “fastboot” feature and the “reload” command after updating to M.10.17, as shown below.

```
ProCurve3400cl#config
ProCurve3400cl(config)# fastboot
ProCurve3400cl(config)# copy tftp flash <ip address of tftp server> M_10_17.swi
The Primary OS Image will be deleted, continue [y/n]? y Validating and Writing System
Software to FLASH...
```

```
ProCurve3400cl(config)# reload
```

```
Device will be rebooted, do you want to continue [y/n]? y
```

```
Rebooting the System
```

Then reconnect and run the show flash command:

```
ProCurve3400cl# show flas
Image           Size(Bytes)   Date   Version
-----
Primary Image   : 3576793    09/26/06 M.10.17
Secondary Image : 3506627    05/26/06 M.10.07
Boot Rom Version: I.08.12
Current Boot    : Primary
```

Please also refer to [“Known Issues” on page 24](#) for additional information regarding updating to the M.10.20 software release.

ProCurve Switch, Routing Switch, and Router Software Keys

Software Letter	ProCurve Networking Products
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. Switch 6108: H.07.xx and earlier
I	Switch 2800 Series (2824 and 2848)
J	Secure Router 7000dl Series (7102dl and 7203dl)
K	Switch 3500yl Series (3500yl-24G-PWR and 3500yl-48G-PWR), Switch 6200yl-24G, 5400zl Series (5406zl, 5406zl-48G, 5412zl, 5412zl-96G), Switch 8212zl and Switch 6600 Series (6600-24G, 6600-24G-4XG, 6600-24XG).
L	Switch 4200vl Series (4204vl, 4208vl, 4202vl-72, and 4202vl-48G)
M	Switch 3400cl Series (3400-24G and 3400-48G): M.08.51 though M.08.97, or M.10.01 and greater; Series 6400cl (6400cl-6XG CX4, and 6410cl-6XG X2): M.08.51 though M.08.95, or M.08.99 to M.08.100 and greater.
N	Switch 2810 Series (2810-24G and 2810-48G)
PA/PB	Switch 1800 Series (Switch 1800-8G – PA.xx; Switch 1800-24G – PB.xx)
Q	Switch 2510 Series (2510-24)
R	Switch 2610 Series (2610-24, 2610-24/12PWR, 2610-24-PWR, 2610-48 and 2610-48-PWR)
T	Switch 2900 Series (2900-24G and 2900-48G)
U	Switch 2510-48
W	Switch 2910al Series (2910al-24G, 2910al-24G-PoE+, 2910al-48G, and 2910al-48G-PoE+)
VA/VB	Switch 1700 Series (Switch 1700-8 - VA and 1700-24 - VB)
WA	ProCurve Access Point 530
WS	ProCurve Wireless Edge Services xl Module and the ProCurve Redundant Wireless Services xl Module
WT	ProCurve Wireless Edge Services zl Module and the ProCurve Redundant Wireless Services zl Module
Y	Switch 2510G Series (2510G-24 and 2510G-48)

Software Management

ProCurve Switch, Routing Switch, and Router Software Keys

Software Letter	ProCurve Networking Products
-----------------	------------------------------

<i>numeric</i>	Switch 9408sl, 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 3400cl 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 www.procurve.com.
2. Click on **Software updates**.
3. Click on **Minimum Software Version Required by Feature**.

For Switch 3400cl Hardware Accessories.

ProCurve Device	Minimum Supported Software Version
J8434A ProCurve 10-GbE Copper Module	M.08.54
J8435A ProCurve 10-GbE Media Flex Module	M.08.54
J8436A ProCurve 10-GbE X2-SC SR Optic	M.08.51
J8437A ProCurve 10-GbE X2-SC LR Optic	M.08.54
J8438A ProCurve 10 GbE X2-SC ER Optic	M.08.75
J8439A ProCurve 10-GbE CX4 Media Converter	M.08.54
J8440A ProCurve 10-GbE X2-CX4 Transceiver	M.08.54
J8440B ProCurve 10-GbE X2-CX4 Transceiver	M.10.06

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, SP2 and 7.0	Sun Java 2 Runtime Environment: – Version 1.5.0_11, Version 1.6.0
Windows Server SE 2003 SP2		
Windows Vista		

Enforcing Switch Security

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

This section provides an overview of switch management and network access security features and applications. For information on specific features, refer to the software manuals provided for your switch model.

Caution:

In its default configuration, the switch is open to unauthorized access of various types. ProCurve recommends that you review this section to help ensure that you recognize the potential for unauthorized switch and network access and are aware of the features available to help prevent such access.

Switch Management Access Security

This section outlines provisions for protecting access to the switch’s status information configuration settings. For more detailed information on these features, refer to the indicated manuals.

Default Settings Affecting Security

In the default configuration, switch management access is available through the following methods:

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

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

Local Manager Password

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

Inbound Telnet Access and Web Browser Access

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

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

(For information on SSH and SSL/TLS, refer to the chapters on these topics in the *Advanced Traffic Management Guide* for your switch.)

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

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

If you choose not to disable Telnet and web browser access, you may want to consider using RADIUS accounting to maintain a record of password-protected access to the switch. Refer to the chapter titled "RADIUS Authentication and Accounting" in the *Access Security Guide* for your switch.

Secure File Transfers

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

SNMP Access (Simple Network Management Protocol)

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

General SNMP Access to the Switch. The switch supports SNMP versions 1, 2c, and 3, including SNMP community and trap configuration. The default configuration supports versions 1 and 2c compatibility, which uses plain text and does not provide security options. ProCurve recommends that you enable SNMP version 3 for improved security. SNMPv3 includes the ability to configure restricted access and to block all non-version 3 messages (which blocks version 1 and 2c unprotected operation). SNMPv3 security options include:

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

For more on SNMPV3, refer to the next subsection and to the chapter titled “Configuring for Network Management Applications” in the *Management and Configuration Guide* for your switch.

SNMP Access to the Switch's Authentication Configuration MIB . A management station running an SNMP networked device management application such as ProCurve Manager Plus (PCM+) or HP OpenView can access the switch's management information base (MIB) for read access to the switch's status and read/write access to the switch's configuration. In earlier software versions, SNMP access to the switch's authentication configuration (hpSwitchAuth) MIB was not allowed. However, beginning with software release M.08.89, the switch's default configuration allows SNMP access to security settings in hpSwitchAuth. If SNMP access to the hpSwitchAuth MIB is considered a security risk in your network, then you should implement the following security precautions when downloading and booting from software release M.08.89 or greater:

1. If SNMP access to the authentication configuration (hpSwitchAuth) MIB described above and in the section titled “[Using SNMP To View and Configure Switch Authentication Features](#)” (page 35) is not desirable for your network, then immediately after downloading and booting from the M.08.89 or greater software for the first time, use the following command to disable this feature:

snmp-server mib hpswitchauthmib excluded

Caution:

Downloading and booting from the M.08.89 or greater software version for the first time enables SNMP access to the authentication configuration MIB (the default action). If SNMPv3 and other security safeguards are not in place, the switch's authentication configuration MIB is exposed to unprotected SNMP access and you should use the above command to disable this access.

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

Refer to “Using SNMP Tools To Manage the Switch” in the chapter titled “Configuring for Network Management Applications” in the Management and Configuration Guide for your switch. .

Physical Access to the Switch

Physical access to the switch allows the following:

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

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

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

For the commands to implement the above actions, refer to “Front-Panel Security” in the chapter titled “Configuring Usernames and Passwords” in the *Access Security Guide* for your switch.

Other Provisions for Management Access Security

Authorized IP Managers. This feature uses IP addresses and masks to determine whether to allow management access to the switch through the network, and covers access through the following:

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

Refer to the chapter titled “Using Authorized IP Managers” in the *Access Security Guide* for your switch.

Secure Management VLAN. This feature creates an isolated network for managing the ProCurve switches that offer this feature. When a secure management VLAN is enabled, CLI, Menu interface, and web browser interface access is restricted to ports configured as members of the VLAN.

Refer to the chapter titled “Static Virtual LANs (VLANs)” in the *Advanced Traffic Management Guide* for your switch.

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

Refer to the chapter titled “RADIUS Authentication and Accounting” in the *Access Security Guide* for your switch.

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

Refer to the chapter titled “TACACS+ Authentication” in the *Access Security Guide* for your switch model.

Access Control Lists (ACLs) for Management Access Protection. ACLs can be used to secure access to the management interface of the switch by blocking inbound IP traffic that has the switch itself as the destination address. (Refer also to “Access Control Lists” in the next section.)

Network Access Security

This section outlines provisions for protecting access through the switch to the network. For more detailed information on these features, refer to the indicated manuals.

Access Control Lists (ACLs)

ACLs enable the switch to permit or deny the following:

- any inbound IP traffic on a port
- specific types of TCP or UDP traffic

While ACLs do not provide user or device authentication, or protection from malicious manipulation of data in IP packet transmissions, ACLs can enhance network security by blocking selected IP traffic types. This functionality can be utilized to:

- permit or deny in-band management access by limiting or preventing the use of designated TCP or UDP protocols
- permit or deny unwanted IP traffic to or from specific hosts

Refer to the chapter titled “Access Control Lists (ACLs) for the Series 3400cl and Series 6400cl Switches” in the *Advanced Traffic Management Guide* for your switch model.

Web and MAC Authentication

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

Refer to the chapter titled “Web and MAC Authentication” in the *Access Security Guide* for your switch model.

Secure Shell (SSH)

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

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

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

Secure Socket Layer (SSLv3/TLSv1)

This feature includes use of Transport Layer Security (TLSv1) to provide remote web access to the switch via authenticated transactions and encrypted paths between the switch and management station clients capable of SSL/TLS operation. The authenticated type includes server certificate authentication with user password authentication.

Refer to the chapter titled “Configuring Secure Socket Layer (SSL) in the *Access Security Guide* for your switch model.

Traffic/Security Filters

These statically configured filters enhance in-band security (and improve control over access to network resources) by forwarding or dropping inbound network traffic according to the configured criteria. Filter options and the devices that support them are listed in the following table:

Switch Model	Source-Port Filters	Protocol Filters	Multicast Filters
Series 6400cl	X	--	--
Series 5400zl	X	X	X
Series 5300xl	X	X	X
Series 4200vl	X	--	--
Series 3500yl	X	X	X
Series 3400cl	X	--	--
Series 2800	X	--	--
Series 2600	X	--	--

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

Refer to the chapter titled “Traffic/Security Filters” in the *Access Security Guide* for your switch model.

802.1X Access Control

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

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

The following table shows the type of access control available on the various ProCurve switch models:

Access Control Types	6200yl 5400zl 3500yl	5300xl 4200vl	3400cl 6400cl	2800 2600 2600-pwr	4100gl
client-based access control (up to 32 authenticated clients per port)	X	X*	--	--	--
port-based access control (one authenticated client opens the port)	X	X	X	X	X
switch operation as a supplicant	X	X	X	X	X
* On the 5300xl switches, this feature is available with software release E.09.02 and greater.					

Refer to the chapter titled “Configuring Port-Based and Client-Based Access Control” in the *Access Security Guide* for your switch model.

Port Security, MAC Lockdown, MAC Lockout, and IP Lockdown

These features provide device-based access security in the following ways:

- **port security:** Enables configuration of each switch port with a unique list of the MAC addresses of devices that are authorized to access the network through that port. This enables individual ports to detect, prevent, and log attempts by unauthorized devices to communicate through the switch. Some switch models also include eavesdrop prevention in the port security feature.
- **MAC lockdown:** This “static addressing” feature is used as an alternative to port security for to prevent station movement and MAC address “hijacking” by allowing a given MAC address to use only one assigned port on the switch. MAC lockdown also restricts the client device to a specific VLAN.
- **MAC lockout:** This feature enables blocking of a specific MAC address so that the switch drops all traffic to or from the specified address.
- **IP lockdown:** Available on Series 2600 and 2800 switches only, this feature enables restriction of incoming traffic on a port to a specific IP address/subnet, and denies all other traffic on that port.

Refer to the chapter titled “Configuring and Monitoring Port Security” in the *Access Security Guide* for your switch model.

Key Management System (KMS)

KMS is available in several ProCurve switch models and is designed to configure and maintain key chains for use with KMS-capable routing protocols that use time-dependent or time-independent keys. (A key chain is a set of keys with a timing mechanism for activating and deactivating individual

keys.) KMS provides specific instances of routing protocols with one or more Send or Accept keys that must be active at the time of a request.

Refer to the chapter titled “Key Management System” in the *Access Security Guide* for your switch model.

Connection-Rate Filtering Based On Virus-Throttling Technology

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

Refer to the chapter titled “Virus Throttling” in the *Access Security Guide* for your switch model.

Identity-Driven Management (IDM)

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

Using IDM, a system administrator can configure automatic and dynamic security to operate at the network edge when a user connects to the network. This operation enables the network to distinguish among different users and what each is authorized to do. Guest access can also be configured without compromising internal security. This means that users can be identified and either approved or denied at the edge of the network instead of in the core.

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

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

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

For more information on IDM, visit the ProCurve web site at <http://www.procurve.com> and click on **Products and Solutions**, then **Identity Driven Management** (under **Network Management**).

Clarifications and Updates

Operating Notes for Jumbo Traffic-Handling

In the Management and Configuration Guide, (Oct., 2005 version) on page 14-33 (page 347 of the .pdf file) where it states:

When a port is not a member of any jumbo-enabled VLAN, it drops all jumbo traffic. If the port is receiving “excessive” inbound jumbo traffic, the port generates an Event Log message to notify you of this condition. This same condition generates a Fault-Finder message in the Alert log of the switch’s web browser interface, and also increments the switch’s “Giant Rx” counter.

Note that it is the “Total Rx Errors” counter that is incremented, not the “Giant Rx” counter. On the 3400cl and 6400cl series switches, when the switch applies the jumbo MTU to a VLAN, all frames with jumbo MTU sizes (1523 to 9220 bytes) are incremented to “Total Rx Errors”.

Non-Genuine Mini-GBIC Detection and Protection Initiative

Non-genuine ProCurve Transceivers and Mini-GBICs have been offered for sale in the marketplace. To protect customer networks from these unsupported products, ProCurve switch software includes the capability to detect and disable non-genuine transceivers and mini-GBICs discovered in Series 3400cl Switch ports. When a non-genuine device is discovered, the switch disables the port and generates an error message in the Event Log.

Publication Updates

Table 1 lists updates to the manual set dated January, 2005.

Table 1. Publication Updates for Manual Set Dated January, 2005

<i>Management and Configuration Guide for the 3400cl, 5300xl, & 6400cl Switches, p/n 5990-6050, January 2005 Edition</i>	Update
Chapter 14: “Configuring for Network Management Applications” Pages 14-44 and 14-49	The show lldp info stats is an invalid command. The correct syntax is: show lldp stats .

IGMP Command Update

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 5990-6051, September 2004 edition. Please refer to this chapter for a detailed explanation of IGMP operation.

The 3400cl switches support the following standards and RFCs:

- 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 3400cl switches:

- Provide full IGMPv2 support as well as full support for IGMPv1 Joins.
- Forward packets for the joined group from all sources, including IGMPv3 Joins.
- Do not support IGMPv3 “Exclude Source” or “Include Source” options in the Join Reports.
- 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.

Using Delayed Group Flush. 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 switches such as the Series 3400cl switches, which support data-driven IGMP. (Data-Driven IGMP, which is enabled by default, prunes off any unregistered 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 (0 - 255 seconds). This command is applied globally to all IGMP-configured VLANs on the switch. A setting of 0 (zero) disables the feature. (Default: Disabled.)

Syntax: `show igmp delayed-flush`

Displays the current setting for the switch.

Setting Fast-Leave and Forced Fast-Leave from the CLI. In earlier switch models, including the 5300xl switches, fast-leave and forced fast-leave options for a port were configured with a lengthy **setmib** command. The following commands now allow a port to be configured for fast-leave or forced fast-leave operation with a conventional CLI command instead of the **setmib** command. Note that these commands must be executed in a VLAN context.

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

*Enables IGMP fast-leaves on the specified ports in the selected VLAN. In the Config context, use the VLAN specifier, for example, **vlan < vid > ip igmp fastleave < port-list >**. The **no** form of the command disables IGMP fast-leave. (Default: Enabled)*

[no] ip igmp forcedfastleave < *port-list* >

Forces IGMP Fast-Leaves on the specified ports in the selected VLAN, even if they are cascaded. (Default: Disabled)

To view a non-default IGMP forced fast-leave configuration on a VLAN, use the **show running-config** command. (The **show running-config** output does not include forced fast-leave if it is set to the default of 0.)

Note

In a future version of the 3400cl switch software, the **show running-config** command output will include any non-default fast-leave settings configured. However, this information is not included in the output for the M.08.53 software release.

IGMP Operating Notes.

- On the Series 3400cl switches, the delayed group flush feature offers little additional benefit over the IGMP data-driven feature (which is enabled by default).
- Forced fast-leave can be used when there are multiple devices attached to a port.

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.

Interoperating with 802.1s Multiple Spanning-Tree

The ProCurve implementation of Multiple Spanning-Tree (MSTP) 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 3400cl 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* for your 3400cl switch.

Rate-Limiting

The configured rate limit on a port reflects the permitted forwarding rate from the port to the switch fabric, 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 “Port Traffic Controls ” of the *Management and Configuration Guide* for your ProCurve Series 3400cl switch. (To download switch documentation, refer to [“Software Updates” on page 1.](#))

Known Issues

Release M.10.17

The following is a known issue related to installation of Release M.10.17 software, which includes a required update to ROM version I.08.12.

When there is an active 10-GbE link in port 26 of the ProCurve 3400cl-24G switch, or port 50 of the ProCurve 3400cl-48G switch, there may be a problem with that link initializing following a software update into the required M.10.17 software version. For customers with a console connection to the switch during the boot process, there may also be a false report with one or more of the following messages:

This switch needs replacement during next scheduled downtime ? or,
Module selftest failure or,
Port [26 or 50] selftest failure. ?

Workarounds: If this is a mission-critical switch and the software is being updated remotely through a 10-GbE link in port 26 or 50, it is recommended that you have someone onsite with the switch able to directly communicate with the switch from another port or the console connection. The issue may be avoided by enabling the “fastboot” feature and using the “reload” command after updating to M.10.17 (refer to [“Install Recommendations for I.08.12 Boot ROM Update” on page 6](#))

If the problem persists, it may also be possible to re-initialize the link by administratively disabling and re-enabling both the affected port and the port that is directly connected to it. If those steps fail to resolve the problem, try disconnecting the media from the potentially affected port until after the switch is running M.10.17. The port should then initialize.

Fix: There is a fix associated with software version M.10.20. Once the switch has been updated to software version M.10.17, update to software version M.10.20 and reboot.

Note that M.10.10 does not have the same issue related to installation of Release M.10.17 software.

Enhancements

Enhancements are listed in chronological order, oldest to newest software release. To review the list of enhancements included since the last general release that was published, begin with [“Release M.10.21 Enhancements” on page 95.](#)

Release M.08.69 Enhancements

Release M.08.69 included the following enhancements:

- Support for Web RADIUS authentication with CLI.
- A new scripting mode.
- Source Port Filter user interface, described in Chapter 9. “Traffic/Security Filters” in the *Access Security Guide* for the switch.

Information on these features is included in the current documentation for the switch, available on the web at: <http://www.hp.com/rnd/support/manuals/>.

Release M.08.70 through M.08.72 Enhancements

Software fixes only; no new enhancements.

Release M.08.73 Enhancements

Release M.08.73 included the following enhancements:

- Support for the new I.08.07 Boot ROM version.
(The 2800/3400/6400 series switches all share the same ROM code)
-

Release M.08.74 through M.08.77 Enhancements

Software fixes only; no new enhancements.

Release M.08.78 Enhancements

Using Fastboot To Reduce Boot Time

The **fastboot** command allows a boot sequence that skips the internal power-on self-tests, resulting in a faster boot time.

Syntax: [no] fastboot

*Used in the global configuration mode to enable the fastboot option. The **no** version of the command disables **fastboot** operation.*

Syntax: show fastboot

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

For example:

```
ProCurve(config)# show fastboot

Fast Boot: Disabled
```

Release M.08.79 Enhancements

CLI Port Rate Display

Beginning with release M.08.79 the CLI “show interface [port list]” command includes the port rate in the display. The rate displayed is the average for a period of 5 minutes, given in bps for 1G ports, or in Kbps for 10G ports. You can also use the CLI command: **show interface port-utilization** to display port-rate over a period of 5 minutes.

The following shows a sample output from this new command.

```
ProCurve# show interface port-utilization
```

Port	Mode	Rx			Tx		
		KBits/s	Pkts/s	Util	KBits/s	Pkts/s	Util
1	100FDx	100000	525	12	100000	400	10
2	1000FDx	0	0	0	0	0	0
3	100FDx	536	44	00.53	504	0	00.50
4	1000FDx	0	0	0	0	0	0
5	1000FDx	0	0	0	0	0	0
6	1000FDx	0	0	0	0	0	0
7	1000FDx	0	5	0	0	0	0
8	1000FDx	0	5	0	0	0	0
9	100FDx	0	30	0	0	0	0

Figure 2. Example rate display output for ports

Operating Notes

- For each port on the switch, the command provides a real-time display of the rate at which data is received (Rx) and transmitted (Tx) in terms of kilobits per second (KBits/s), number of packets per second (Pkts/s), and utilization (Util) expressed as a percentage of the total bandwidth available.
- As in previous software versions, the **show interfaces** <port-list> command can be used to display the current link status and the port rate average over a 5 minute period. Port rates are shown in bits per second (bps) for ports up to 1 Gigabit, and are shown in kilobits per second (Kbps) for 10 Gigabit ports.

Release M.08.80 through M.08.83 Enhancements

Software fixes only; no new enhancements.

Release M.08.84 Enhancements

Release M.08.84 includes the following enhancement:

Added the `show tech transceivers` command to allow removable transceiver serial numbers to be read without removal of the transceivers from the switch. :

Release M.08.85 through M.08.88 Enhancements

Software fixes only; no new enhancements.

Release M.08.89 Enhancements

Release M.08.89 includes the following enhancements:

- DNS Resolver for using DNS names for Ping and Traceroute
- RADIUS Configuration via SNMP (see [“Using SNMP To View and Configure Switch Authentication Features” on page 35](#))

DNS Resolver

The Domain Name System (DNS) resolver is designed for use in local network domains where it enables use of a host name or fully qualified domain name to perform **ping** and **traceroute** operations from the switch.

Terminology

Domain Suffix — Includes all labels to the right of the unique host name in a fully qualified domain name assigned to an IP address. For example, in the fully qualified domain name “device53.evergreen.trees.org”, the domain suffix is “evergreen.trees.org”, while “device53” is the unique (host) name assigned to a specific IP address.

Fully Qualified Domain Name — The sequence of labels in a domain name identifying a specific host (host name) and the domain in which it exists. For example, if a device with an IP address of 10.10.10.101 has a host name of *device53* and resides in the *evergreen.trees.org* domain, then the device’s fully qualified domain name is *device53.evergreen.trees.org* and the DNS resolution of this name is 10.10.10.101.

Host Name — The unique, leftmost label in a domain name assigned to a specific IP address in a DNS server configuration. This enables the server to distinguish a device using that IP address from other devices in the same domain. For example, in the *evergreen.trees.org* domain, if an

IP address of 10.10.100.27 is assigned a host name of *accounts015* and another IP address of 10.10.100.33 is assigned a host name of *sales021*, then the switch configured with the domain suffix *evergreen.trees.org* and a DNS server that resolves addresses in that domain can use the host names to reach the devices with **ping** and **traceroute** commands:

```
ping accounts015
traceroute sales021
```

Basic Operation

- When the switch is configured with only the IP address of a DNS server available to the switch, then a **ping** or **traceroute** command, executed with a fully qualified domain name, can reach a device found in any domain accessible through the configured DNS server.

- When the switch is configured with both of the following:
 - the IP address of a DNS server available to the switch
 - the domain suffix of a domain available to the configured DNS server

then:

- A **ping** or **traceroute** command that includes the host name of a device in the same domain as the configured domain suffix can reach that device.
- A **ping** or **traceroute** command that includes a fully qualified domain name can reach a device in any domain that is available to the configured DNS server.

Example. Suppose the switch is configured with the domain suffix **mygroup.procurve.net** and the IP address for an accessible DNS server. If an operator wants to use the switch to ping a host using the DNS name “leader” assigned to an IP address used in that domain, then the operator can use either of the following commands:




Figure 3. Example of Using Either a Host Name or a Fully Qualified Domain Name

In the preceding example, if the DNS server’s IP address is configured on the switch, but a domain suffix is not configured, then the fully qualified domain name *must* be used.

Note that if the target host is in a domain *other than* the domain configured on the switch, then:

- The host's domain must be reachable from the switch. This requires that the DNS server for the switch must be able to communicate with the DNS server(s) in the path to the domain in which the target host operates.
- The fully qualified domain name must be used, and the domain suffix must correspond to the domain in which the target host operates, regardless of the domain suffix configured in the switch.

Example. Suppose the switch is configured with the domain suffix **mygroup.procurve.net** and the IP address for an accessible DNS server. This time, the operator wants to use the switch to trace the route to a host named "remote-01" in another domain named **common.group.net**. As long as this domain is accessible to the DNS server configured on the switch, a **traceroute** command using the target's fully qualified DNS name should succeed.

```
ProCurve# traceroute [remote-01.common.group.net] ← Fully Qualified Host Name for the Target Host
[traceroute to 10.22.240.73] ← IP Address for Target Host "remote-01"
      1 hop min, 30 hops max, 5 sec. timeout, 3 probes
 1 10.28.229.3          0 ms          0 ms          0 ms
 2 10.71.217.1         0 ms          0 ms          0 ms
 3 10.0.198.2          1 ms          0 ms          0 ms
[4 10.22.240.73       0 ms          0 ms          0 ms]
[ _ _ _ _ _ ]
```

Figure 4. Example Using the Fully Qualified Domain Name for an Accessible Target in Another Domain

Configuring and Using DNS Resolution with Ping and Traceroute Commands

1. Determine the following:
 - a. the IP address for a DNS server operating in a domain in your network
 - b. the domain name for an accessible domain in which there are hosts you want to reach with **ping** and/or **traceroute** commands. (This is the domain suffix in the fully qualified domain name for a given host operating in the selected domain. Refer to [“Terminology” on page 28.](#)) Note that if a domain suffix is not configured, fully qualified domain names can be used to resolve **ping** and **traceroute** commands.
 - c. the host names assigned to target IP addresses in the DNS server for the specified domain
2. Use the data from steps 1a and 1b to configure the DNS entry on the switch.
3. Use either **ping** or **traceroute** with the host names for the target devices whose connectivity you are testing or troubleshooting.

Configuring a DNS Entry

The switch allows one DNS server entry, which includes the DNS server IP address and the chosen domain name suffix. Configuring the entry enables the use of **ping** and **traceroute** with a target's host name instead of the target's IP address.

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

*Configures the IP address of a DNS server accessible to the switch. This setting identifies the server to use for DNS resolution to the target IP address, and must be configured before **ping** or **traceroute** can be executed with host name criteria.*

The switch supports one DNS server entry. Configuring another IP address for this value replaces the current IP address with the new one.

*The **no** form of the command replaces the configured IP address with the null setting, which disables host name resolution. (Default: null)*

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

*Configures the domain suffix that is automatically appended to the host name entered with the **ping** or **traceroute** command. When the domain suffix and the DNS server IP address are both configured on the switch, you can execute **ping** and **traceroute** with only the host name of the desired target within the domain. In either of the following two instances, you must manually provide the domain identification by using a fully qualified DNS name with each **ping** and **traceroute** command:*

- *If the DNS server IP address is configured on the switch, but the domain suffix is not configured (null)*
- *The domain suffix configured on the switch is not the domain in which the target host exists*

The switch supports one domain suffix entry. Configuring a new entry for this value replaces the current suffix.

*The **no** form of the command replaces the configured domain suffix with the null setting. (Default: null)*

Example Using DNS Names with Ping and Traceroute

In the network illustrated in [Figure 5](#), the switch at 10.28.192.1 is configured to use DNS names for **ping** and **traceroute** in the *pubs.outdoors.com* domain. The DNS server has been configured to assign the host name *docservr* to the IP address used by the document server (10.28.229.219).




Figure 5. Example Network Domain

Configuring switch "A" with the domain name and the IP address of a DNS server for the domain enables the switch to use host names assigned to IP addresses in the domain to perform **ping** and **traceroute** actions on the devices in the domain. To summarize:

Entity:	Identity:
DNS Server IP Address	10.28.229.10
Domain Name (and Domain Suffix for Hosts in the Domain)	pubs.outdoors.com
Host Name Assigned to 10.28.229.219 by the DNS Server	docservr
Fully Qualified Domain Name for the IP address Used By the Document Server (10.28.229.219)	docservr.pubs.outdoors.com
Switch IP Address	10.28.192.1
Document Server IP Address	10.28.229.219

With the above already configured, the following commands enable **ping** and **traceroute** with the host name **docservr** to reach the document server at 10.28.229.219.

```
ProCurve(config)# ip dns server-address 10.28.229.10
ProCurve(config)# ip dns domain-name pubs.outdoors.com
```

Figure 6. Configuring Switch "A" in Figure 5 To Support DNS Resolution

```
ProCurve# ping docservr
10.28.229.219 is alive, time = 1 ms

ProCurve# traceroute docservr
traceroute to 10.28.229.219
      1 hop min, 30 hops max, 5 sec. timeout, 3 probes
 1 10.28.192.2          1 ms      0 ms      0 ms
 2 10.28.229.219       0 ms      0 ms      0 ms
```

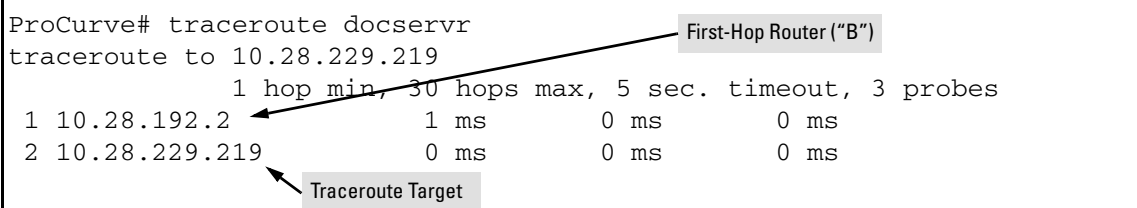


Figure 7. Example of Ping and Traceroute Execution for the Network in Figure 5 on Page 32

As mentioned under “Basic Operation” on page 29, if the DNS entry configured in the switch includes only the DNS server’s IP address, you must use the target host’s fully qualified domain name with **ping** and **traceroute**. For example, using the document server in Figure 5 as a target:

```
ProCurve# ping docservr.pubs.outdoors.com
10.28.229.219 is alive, time = 1 ms

ProCurve# traceroute docservr.pubs.outdoors.com
traceroute to 10.28.229.219
      1 hop min, 30 hops max, 5 sec. timeout, 3 probes
 1 10.28.192.2          1 ms      0 ms      0 ms
 2 10.28.229.219       0 ms      0 ms      0 ms
```

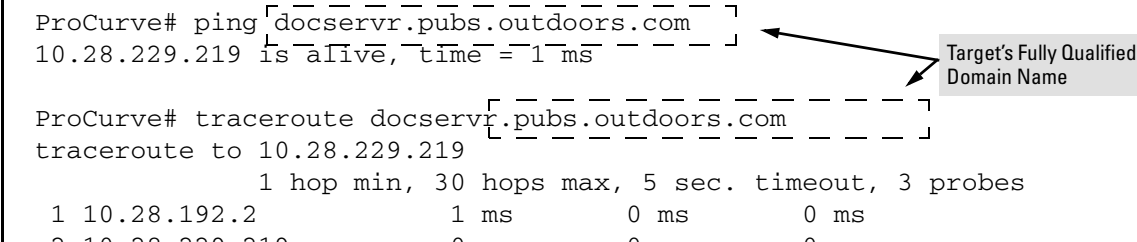


Figure 8. Example of Ping and Traceroute Execution When Only the DNS Server IP Address Is Configured

Viewing the Current DNS Configuration

The **show ip** command displays the current DNS configuration along with other IP configuration information. If the switch configuration currently includes a nondefault (non-null) DNS entry, it will also appear in the **show run** command output.

```
ProCurve# show ip

Internet (IP) Service

IP Routing : Disabled

Default Gateway : 10.28.192.2
Default TTL     : 64
Arp Age        : 20
Domain Suffix   : pubs.outdoors.com
DNS server      : 10.28.229.10

VLAN          | IP Config  IP Address      Subnet Mask
-----+-----
DEFAULT_VLAN | Manual     10.28.192.1     255.255.255.0
```




Figure 9. Example of Viewing the Current DNS Configuration

Operating Notes

- The DNS server must be accessible to the switch, but it is not necessary for any intermediate devices between the switch and the DNS server to be configured to support DNS operation.
- A DNS configuration must include the IP address for a DNS server that is able to resolve host names for the desired domain. If a DNS server has limited knowledge of other domains, then its ability to resolve **ping** or **traceroute** requests is also limited.
- If the DNS configuration includes a DNS server IP address but does not also include a domain suffix, then any ping or traceroute command should include the target host's fully qualified domain name. Refer to [Figure 3](#) on page 29.
- The switch supports one DNS entry; that is, one DNS server IP address and the corresponding domain name suffix.
- Switch-Initiated DNS packets go out through the VLAN having the best route to the DNS server, even if a Management VLAN has been configured.
- The **traceroute** command output shows only IP addresses.
- The DNS server address must be manually input. It is not be automatically determined via DHCP.
- Operation with IPv4 DNS servers has been verified and, while no problems with servers supporting both IPv4 and IPv6 addresses are expected, testing has not been performed with such servers. (IPv6 AAAA-style queries are not supported.)

Event Log Messages

Message	Meaning
DNS server address not configured	The switch does not have an IP address configured for the DNS server.
DNS server not responding	The DNS server failed to respond or is unreachable. An incorrect server IP address can produce this result.
Unknown host < <i>host-name</i> >	The host name did not resolve to an IP address. Some reasons for this occurring include: <ul style="list-style-type: none">• The host name was not found.• The named domain was not found.• The domain suffix was expected, but has not been configured. (If the server's IP address has been configured in the switch but the domain name has not been configured, then the host's fully qualified domain name must be used.)

Using SNMP To View and Configure Switch Authentication Features

In earlier software releases, SNMP MIB object access has not been available for switch authentication configuration (hpSwitchAuth) features. Beginning with software release M.08.89, the 3400cl and 6400cl switches allow, by default, manager-only SNMP read/write access to a subset of the authentication MIB objects for the following features:

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

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

Security Notes

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

To help prevent unauthorized access to the switch's authentication MIB, ProCurve recommends enhancing security according to the guidelines under [“Enforcing Switch Security” on page 10](#).

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

If you choose to leave SNMP access to the security MIB open (the default setting), ProCurve recommends that you configure the switch with the SNMP version 3 management and access security feature, and disable SNMP version 2c access. (Refer to [“Enforcing Switch Security” on page 10](#).)

Changing and Viewing the SNMP Access Configuration

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

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

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

(Default: included)

Syntax: show snmp-server

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

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

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

SNMP Communities

Community Name	MIB View	Write	Access
public	Manager	Unrestricted	

Trap Receivers

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

Send Authentication Traps [No] : No

Address	Community	Events Sent in Trap

Excluded MIBs

hpSwitchAuthenticationMIB

This command disables SNMP security MIB access.

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

Figure 10. Disabling SNMP Access to the Authentication MIB and Displaying the Result

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

```
ProCurve(config)# show run

Running configuration:

; J4905A Configuration Editor; Created on release #M.10.05

hostname "ProCurve"
[snmp-server mib hpSwitchAuthMIB excluded ] ←
ip default-gateway 10.10.24.55
snmp-server community "public" Operator
vlan 1
    name "DEFAULT_VLAN"
    untagged 1-26
    ip address 10.10.24.100 255.255.255.0
    exit
password manager
```

Indicates that SNMP access to the authentication configuration MIB (hpSwitchAuth) is disabled.

Figure 11. Using the show run Command to View the Current Authentication MIB Access State

Releases M.08.90 and M.08.91 Enhancements

- The MSTP enhancement implementing the CLI command for spanning-tree legacy-path-cost was included in release M.08.90
- The MSTP enhancement implementing the CLI command for spanning-tree legacy-mode was included in release M.08.91
- QoS Pass-Through Mode enhancement, a new command that allows the configuration of the Quality of Service (QoS) queues to be selected.

MSTP Default Path Cost Controls

Summary: 802.1D and 802.1t specify different default path-cost values (based on interface speed). These are used if the user hasn't configured a "custom" path-cost for the interface. The default of this toggle is to use 802.1t values. The reason one might set this control to 802.1D would be for better interoperability with legacy 802.1D STP (Spanning Tree Protocol) bridges.

To support legacy STP bridges, the following commands (options) have been added to CLI:

- spanning-tree legacy-path-cost** - Use 802.1D values for default path-cost
- no spanning-tree legacy-path-cost** - Use 802.1t values for default path-cost (default setting)

The “legacy-path-cost” CLI command does not affect or replace functionality of the “spanning-tree force-version” command. The “spanning-tree force-version” controls whether MSTP will send and process 802.1w RSTP, or 802.1D STP BPDUs. Regardless of what the “legacy-path-cost” parameter is set to, MSTP will interoperate with legacy STP bridges (send/receive Config and TCN BPDUs).

spanning-tree legacy-mode - A “macro” that is the equivalent of executing the “spanning-tree legacy-path-cost” and “spanning-tree force-version stp-compatible” commands.

no spanning-tree legacy-mode - A “macro” that is the equivalent of executing the “no spanning-tree legacy-path-cost” and “spanning-tree force-version mstp-compatible” commands.

When either legacy-mode or legacy-path-cost control is toggled, all default path costs will be recalculated to correspond to the new setting, and spanning tree is recalculated if needed.

QoS Pass-Through Mode

Release M.08.91 introduced a new command that allows the configuration of the Quality of Service (QoS) queues to be selected. By better matching the configuration of the QoS queues to the amount of prioritized and non-prioritized traffic being transferred, performance can be improved and packet loss due to over-subscription can be minimized.

In previous software versions, the 3400cl and the 6400cl switches had four QoS queues of equal size. Depending on the mix of prioritized and non-prioritized traffic, this configuration might not always optimize performance and could result in dropped packets when resources were over-subscribed. Starting with this software version, four QoS Pass-Through modes are available for use. The number of queues and the size of the memory buffer used by each queue differs in each mode. [Table 2](#) below summarizes the QoS queue configuration of each mode

Table 2. QoS Pass-Through Modes

QoS Pass-Through Mode	Number of Queues	QoS Queue Memory Buffer Configuration	Description
typical (default)	4	One large queue for Priority 0 and 3 traffic and three other queues for the remaining traffic.	A mix of prioritized and non-prioritized traffic. This is the default mode, used when QoS Pass-Through is disabled.
balanced	4	All queues are the same size.	Equal amounts of prioritized and non-prioritized traffic. This is the same mode used in pre-M.08.78 software versions.
one-queue	1	One large queue. ¹	No traffic is prioritized.
optimized	2	One small queue for Priority 6 and 7 traffic; one large queue for all other traffic.	Most traffic is not prioritized.

¹This mode has a small queue used exclusively for Priority 7 management and control traffic.

Note

Changing the QoS Pass-Through Mode can be done without rebooting the switch. However, the switch ports are toggled down and back up, allowing the QoS queues to be reconfigured. This may affect routing and spanning tree operation. ProCurve Networking recommends that QoS queues be reconfigured during periods of non-peak traffic.

Configuring QoS Pass-Through Mode

Syntax: qos-passthrough-mode [balanced | one-queue | optimized | typical]

*Specifies the QoS queue mode to be used by the switch. The number of queues and the size of each queue is determined by the mode selected. If no mode is specified the **optimized** mode is used. QoS Pass-Through is disabled using the **no qos-pass-through** command.*

balanced: *Configures four QoS queues of the same size. This configuration is the same as was used by software versions prior to M.08.78.*

one-queue: *Configures one QoS queue. By consolidating packet buffer memory, line-rate flows with no loss of data may be achieved.*

Note: *This mode has a small queue used exclusively by Priority 7 management and control packets.*

optimized: *Configures two QoS queues: a small queue for Priority 6 and 7 traffic and a large queue for all other traffic.*

typical: *Configures four QoS queues: a large queue for Priority 0 and 3 traffic, and three other queues for the remaining traffic. This is the default configuration on the switch and is used when QoS Pass-Through is disabled.*

Syntax: [no] qos-passthrough-mode

*Specifies the **optimized** QoS queue mode for the switch.*

*The **no qos-pass-through** command returns the QoS queue mode to **typical**, the default setting.*

Configuring QoS Pass-Through Mode Through the CLI. The following example changes the QoS Pass-Through Mode to **balanced**. A **show** command verifies the new mode.

```
ProCurve(config)# qos-passthrough-mode balanced
This requires a temporary shut-down of logical ports. Continue (y/n) y ←
ProCurve(config)# show qos-passthrough-mode

Qos passthrough mode : balanced

ProCurve(config)#
```

Reconfiguring the QoS queues toggles the switch ports, which may affect routing and spanning tree operation. Choose **n** to cancel this operation.

Figure 1. Example Showing QoS Pass-Through Mode Set Using the CLI

QoS Pass-Through Mode SNMP MIB Object. A read-write MIB object, 1.3.6.1.4.1.11.2.14.11.5.1.7.1.24.1, has been added to the ProCurve switch MIB. The QoS Pass-Through Mode can be changed using either an SNMP network management application or the CLI **setmib** command.

Syntax: setMIB hpSwitchQosPassThroughModeConfig.0 -i [1 | 2 | 3 | 4]

Specifies the QoS queue mode to be used by the switch. The number of queues and the size of each queue is determined by the mode selected.

- 1 optimized:** Configures two QoS queues: a small queue for Priority 6 and 7 traffic and a large queue for all other traffic.
- 2 typical:** Configures four QoS queues: a large queue for Priority 0 and 3 traffic, and three other queues for the remaining traffic. This is the default configuration on the switch and is used when QoS Pass-Through is disabled.
- 3 balanced:** Configures four QoS queues of the same size. This configuration is the same as was used by software versions prior to M.08.xx.
- 4 one-queue:** Configures one QoS queue. By consolidating packet buffer memory, line-rate flows with no loss of data may be achieved.
Note: This mode has a small queue used exclusively by Priority 7 management and control packets.

The following example changes the QoS Pass-Through Mode to **one-queue**. A **show** command verifies the new mode.

```
ProCurve(config)# setMIB hpSwitchQosPassThroughModeConfig.0 -i 4
hpSwitchQosPassThroughModeConfig.0 = 4
ProCurve(config)# show qos-passthrough-mode

Qos passthrough mode : one-queue

ProCurve(config)#
```

Figure 2. Example Showing QoS Pass-Through Mode Set Using the setMIB Command

Displaying the Current QoS Pass-Through Mode on the Switch

The following command indicates the current QoS Pass-Through Mode on the switch.

Syntax: show qos-passthrough-mode

*This command displays the current QoS Pass-Through Mode configured on the switch. The default mode is **typical**.*

The current QoS Pass-Through Mode also is displayed in the **show running-config** command output.

Operating Notes

- To use the same QoS queue structure used in pre-M.08.78 software, set the QoS Pass-Through Mode to **balanced**.
- The **optimized** mode matches the QoS Pass-through mode on the ProCurve Series 2800 switches. This mode is used when the QoS Pass-Through Mode command is entered with no arguments, **qos-passthrough-mode**.

Release M.08.94 Enhancements

Release M.08.94 includes the following enhancements:

- Added DHCP Option 82 functionality for 3400cl series.
- UDP broadcast forwarding feature is now supported on the 3400cl series.

DHCP Option 82: Using the Management VLAN IP Address for the Remote ID

This section describes the Management VLAN enhancement to the DHCP option 82 feature. For more information on DHCP option 82 operation, refer to “Configuring DHCP Relay” in the chapter titled “IP Routing Features” in the *Advanced Traffic Management Guide*.

When the routing switch is used as a DHCP relay agent with Option 82 enabled, it inserts a relay agent information option into client-originated DHCP packets being forwarded to a DHCP server. The option automatically includes two suboptions:

- Circuit ID: the identity of the port through which the DHCP request entered the relay agent
- Remote ID: the identity (IP address) of the DHCP relay agent

Using earlier software releases, the remote ID can be either the routing switch’s MAC address (the default option) or the IP address of the VLAN or subnet on which the client DHCP request was received. Beginning with software release M.08.xx, if a Management VLAN is configured on the routing switch, then the Management VLAN IP address can be used as the remote ID.

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

[ip | mac | mgmt-vlan] : Specifies the remote ID suboption the routing switch will use in Option 82 fields added or appended to DHCP client packets. The choice depends on how you want to define DHCP policy areas in the client requests sent to the DHCP server. If a remote ID suboption is not configured, then the routing switch defaults to the **mac** option.

mgmt-vlan: Specifies the IP address of the (optional) Management VLAN configured on the routing switch. Requires that a Management VLAN is already configured on the switch. If the Management VLAN is multinetted, then the primary IP address configured for the Management VLAN is used for the remote ID.

ip: Specifies the IP address of the VLAN on which the client DHCP packet enters the routing switch. In the case of a multinetted VLAN, the remote ID suboption uses the IP address of the subnet on which the client request packet is received.

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.)
(Default: **mac**)

Example

In the routing switch shown below, option 82 has been configured with **mgmt-vlan** for the Remote ID.

```
ProCurve(config)# dhcp-relay option 82 append mgmt-vlan
```

The resulting effect on DHCP operation for clients X, Y, and Z is shown in [Table 3](#).

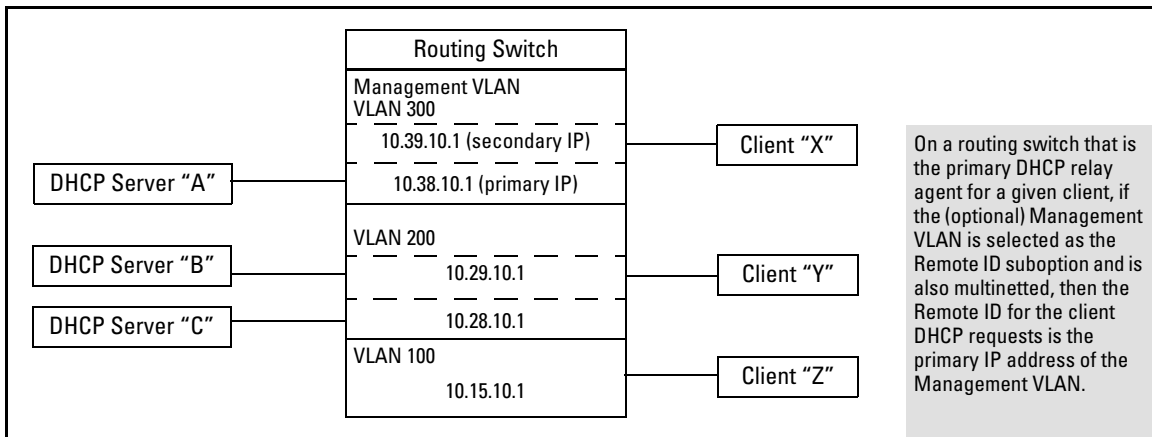


Figure 12. DHCP Option 82 When Using the Management VLAN as the Remote ID Suboption

Table 3. DHCP Operation for the Topology in Figure 12

Client	Remote ID	giaddr*	DHCP Server	
X	10.38.10.1	10.39.10.1	A only	If a DHCP client is in the Management VLAN, then its DHCP requests can go only to a DHCP server that is also in the Management VLAN. Routing to other VLANs is not allowed.
Y	10.38.10.1	10.29.10.1	B or C	Clients outside of the Management VLAN can send DHCP requests only to DHCP servers outside of the Management VLAN. Routing to the Management VLAN is not allowed.
Z	10.38.10.1	10.15.10.1	B or C	

*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*). This 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.

Operating Notes

- Routing is not allowed between the Management VLAN and other VLANs. Thus, a DHCP server must be available in the Management VLAN if there are clients in the Management VLAN that require a DHCP server.
- If the Management VLAN IP address configuration changes after **mgmt-vlan** has been configured as the remote ID suboption, the routing switch dynamically adjusts to the new IP addressing for all future DHCP requests.
- The Management VLAN and all other VLANs on the routing switch use the same MAC address.

UDP Broadcast Forwarding

Beginning with software release M.08.94, UDP Broadcast Forwarding is available on the ProCurve 3400cl and 6400cl switches. For further information, refer to the section titled “UDP Broadcast Forwarding on 5300xl Switches” in the “IP Routing Features” chapter of the *Advanced Traffic Management Guide* for your switch. (Note that this manual covers multiple switches and the description of UDP Broadcast Forwarding is no longer restricted to just the 5300xl switches.)

Some applications rely on client requests sent as limited IP broadcasts addressed to a UDP application port. If a server for the application receives such a broadcast, the server can reply to the client. Since typical router behavior, by default, does not allow broadcast forwarding, a client’s UDP broadcast requests cannot reach a target server on a different subnet unless the router is configured to forward client UDP broadcasts to that server.

A switch with routing enabled includes optional per-VLAN UDP broadcast forwarding that allows up to 256 server and/or subnet entries on the switch (16 entries per-VLAN). If an entry for a particular UDP port number is configured on a VLAN and an inbound UDP broadcast packet with that port number is received on the VLAN, then the switch routes the packet to the appropriate subnet. (Each entry can designate either a single device or a single subnet. The switch ignores any entry that designates multiple subnets.)

Releases M.08.95 through M.10.01 Enhancements

Software fixes only; no new enhancements.

Release M.08.96 Enhancements

- Enabled use of login "Message of the Day" (MOTD) banner. For details on using this feature, refer to "Custom Login Banners for the Console and Web Browser Interfaces" in Chapter 2 of the *Management and Configuration Guide* for 3400cl and 6400cl switches.
-

Releases M.08.97 through M.10.01 Enhancements

No new enhancements in release M.08.97. The M code software for the 3400cl then branched to M.10.01, which has software fixes only, no enhancements.

Release M.10.02 Enhancements

Release M.10.02 includes the following enhancements:

- Support for RADIUS assigned ACLs (access control lists).
- Added new "show sFlow" commands.

RADIUS-Assigned Access Control Lists (ACLs)

Introduced with software release M.10.*xx* on the 3400cl switches, 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. The ACL is applied to the switch port used by the client and remains in force for the duration of the client session. 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 for a given client requires the following:

- RADIUS authentication of the client must be available on the switch through either 802.1X, Web authentication, or MAC authentication.
-

- An ACL must be configured on the RADIUS server (instead of the switch) by creating and assigning one or more Access Control Entries to the username/password pair or MAC address of the client for which you want ACL support.
- Where 802.1X is used for client authentication, then either the client device must be running 802.1X supplicant software or the capability must exist for the client to download this software from the network through use of the 802.1X Open VLAN mode available on the switch. (If authentication is achieved through Web or MAC Authentication, then 802.1X supplicant software is not required.)

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 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.

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 4](#), highlights several key differences between the static ACLs configurable on 3400cl switch ports and the dynamic ACLs that can be assigned to individual ports by a RADIUS server. (The switch supports either one RADIUS-based ACL or one port-based ACL at a time on a given port. It does not support having both ACL types on the same port at the same time.)

Table 4. Contrasting Dynamic and Static ACLs

RADIUS-Based (Dynamic) ACLs	Port-Based (Static) ACLs
Operates on the 3400cl switches.	Operates on both the 3400cl and 6400cl switches.
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 the traffic to the switch from connected devices is predictable and largely static.
Implementation requires client authentication.	Client authentication not a factor.
Instead of an ACL name or number, the ACL is defined by the credentials (username/password pair or the MAC address) of the specific client the ACL is intended to service. Thus, all ACEs configured in the RADIUS server with the same client identifiers comprise the ACL for the specified client.	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 traffic from a connected device, and operates in applications that may or may not include 802.1X or other types of client authentication.
When the authenticated client session ends, the switch removes the RADIUS-assigned ACL from the client port.	Remains statically assigned to the ports unless removed by a no interface < port-list > access-group CLI command.
Supports one RADIUS-based ACL on a port.	Supports one inbound ACL per-port.
The ACL filters the IP traffic received inbound from the client whose authentication resulted in the ACL assignment. Inbound traffic from any other source is denied.	An ACL applied inbound on a port filters all IP traffic received.
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 all inbound IP traffic on the specified ports.
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)
- [*cnt*] (optional counter that increments when there is a packet match)

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.

Extended ACL: This type of Access Control List uses layer-3 IP criteria composed of source and destination IP addresses and (optionally) TCP or UDP port criteria to determine whether there is a match with an IP packet. On the 3400cl switches, the source IP address is always defined as “any”, and extended ACLs apply only to inbound bridged or routed traffic. For a RADIUS-based, extended ACL assigned to a port, only the inbound traffic from the client whose authentication caused the ACL assignment is filtered. Inbound traffic from any other sources is denied.

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 (from the authenticated client) 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 the authenticated source* to any destination. Inbound traffic from any other sources is denied. (Inbound traffic from a client *other than* the client whose authentication caused in the ACL assignment will be denied.)

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 48.

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 caused a RADIUS-based ACL assignment. Inbound traffic from any other source, including a second, authenticated client (on the same port) will be denied.

The Packet-filtering Process

Sequential Comparison and Action. When an ACL filters a packet from an authenticated client, 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 from the authenticated client 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 (from the authenticated client) 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". (Note that the "permit any" option applies only to packets from the client whose authentication caused the assignment of the ACL to the port.)

Example. Suppose the ACL in [Figure 3](#) is assigned to filter the traffic from an authenticated client on a given port in the switch:

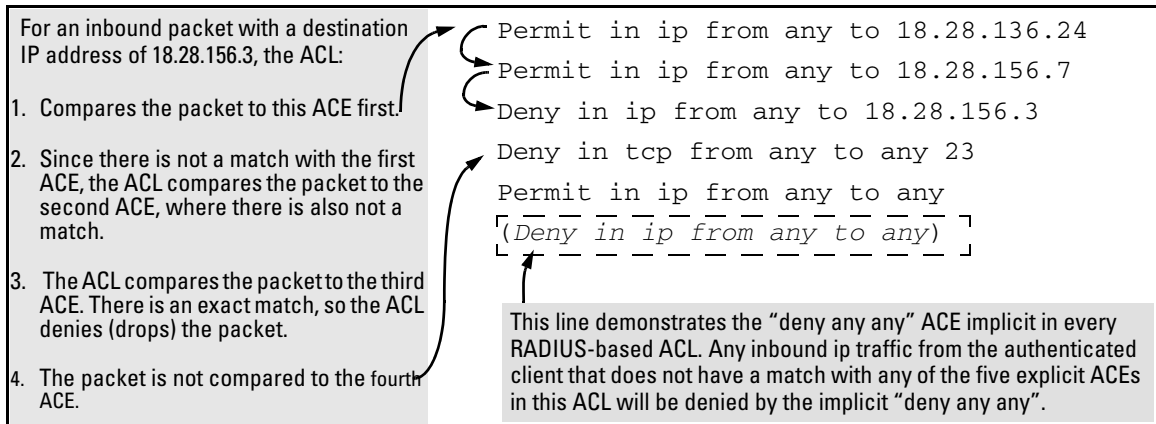



Figure 3. Example of Sequential Comparison

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

Because only one ACL is allowed on a port, if a statically configured ACL already exists on a port, a RADIUS-based ACL cannot be assigned to that port. In this case, if a client authenticates and the RADIUS server is configured to assign a dynamic ACL to the port for that client, the client will then be de-authenticated.



1. If a match is not found with the first ACE in an ACL, the switch proceeds to the next ACE and so on.
2. If a match with an explicit ACE is subsequently found, the packet is either permitted (forwarded) or denied (dropped), depending on the action specified in the matching ACE. In this case the switch ignores all subsequent ACEs in the ACL.
3. If a match is not found with any explicit ACE in the ACL, the switch invokes the implicit **deny IP any** at the end of every ACL, and drops the packet.

Note: If the list includes a **permit IP any** entry, no packets can reach the implicit **deny IP any** at the end of the list. Also, a **permit IP any** ACE at any point in an ACL defeats the purpose of any subsequent ACEs in the list.

Figure 4. 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 from the authenticated client, 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:

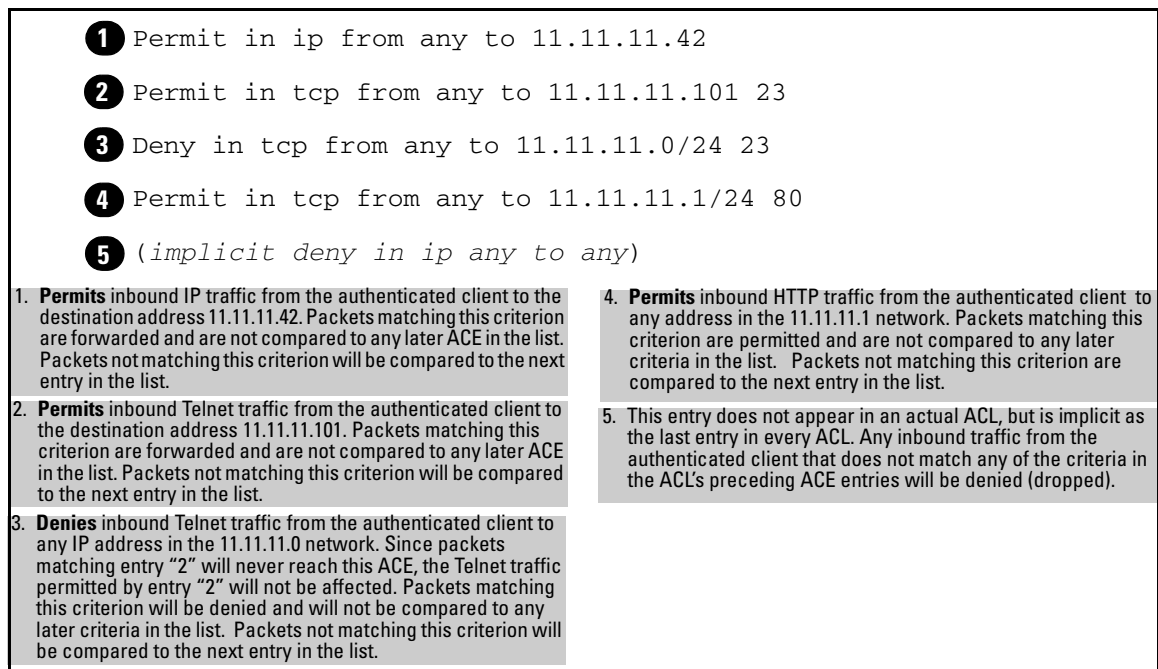


Figure 5. Example of How a RADIUS-Based ACL Filters Packets

Overriding the Implicit “deny IP any any”. 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. (Inbound traffic from a client other than the client whose authentication caused the ACL assignment to the port is dropped.)

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 policies 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? In some cases you will need to explicitly identify permitted traffic. In other cases, depending on your policies, you can insert a **permit in ip from any to any** entry at the end of the ACL so that all IP traffic (from the authenticated client) that is 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 in ip from any to 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 5 on page 57.](#))

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 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 port, 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/Package 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 from the Authenticated Client:** Entering a **permit in ip from any to any** (permit any any) ACE in a RADIUS-based ACL permits all IP traffic (from the authenticated client) that is not previously permitted or denied by that ACL. Any ACEs listed after that point do not have any effect. (While a RADIUS-based ACL is applied to a port, traffic inbound from sources other than the client whose authentication caused the ACL assignment is denied.)

- **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 5](#) on page 53.
- 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

- **ACL Assignments Per-Port:** One RADIUS-assigned ACL is allowed per-port.
- **Port Trunks Excluded:** RADIUS-assigned ACLs cannot be assigned to a port trunk.
- **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 58.
- **Multiple Clients Using the Same Username/Password Pair:** Multiple clients using the same username\password pair will use duplicate instances of the same ACL.
- **RADIUS-Based ACL Not Allowed on a Port that has a Statically-Configured ACL:** Where a RADIUS server is configured to assign an ACL when a given client authenticates, if the port used by that client is already statically configured with a port-based ACL in the switch configuration, then the RADIUS-based ACL is not accepted and the client is de-authenticated.
- **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.

Limits for RADIUS-Based ACLs, Associated ACEs, and Counters

Table 5 describes limits the switch supports in ACLs applied by a RADIUS server. Exceeding a limit causes the related client authentication to fail.

Table 5. Limits Affecting RADIUS-Based ACL Applications

Item	Limit	Notes	
Maximum Number of Authenticated Client Sessions Per-Port Using RADIUS-based ACLs	1	One RADIUS-based ACL can operate on a given port at a time. If an authenticated client is already using a RADIUS-based ACL on a port and a second client requiring a RADIUS-based ACL attempts to authenticate on the same port, the attempt by the second client will fail.	
Maximum Number of (internal) ACEs Per-Port, and Maximum Number of (internal) ACEs Per-ACL	Up to 120*	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
*Uses shared internal resources, which can affect the per-port availability of internal ACEs. Refer to the section titled "Planning an ACL Application on a Series 3400cl or 6400cl Switch" in the chapter titled "Access Control Lists (ACLs) for the Series 3400cl and 6400cl Switches" in the <i>Advanced Traffic Management Guide</i> for your switch model. Use the show access-list resources command to view the current resources available for the ports on the switch.			
Maximum Number of Characters in an ACE	80	—	
Maximum Number of (optional) Internal Counters Used Per-ACL	32	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 ACE examples below show how the switch calculates internal counter groups.	
		Examples of ACE Usage of Internal Counters	
		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

Item	Limit	Notes
Per-Port Mask Usage		ACLs consume per-port (internal) mask resources rapidly and can be affected by IGMP usage on the same switch. For more on this topic, refer to the “ACL Resource Usage and Monitoring” and “Extended ACLs” subsections in the chapter titled “Access Control Lists (ACLs) for the Series 3400cl and Series 6400cl Switches” of the <i>Advanced Traffic Management Guide</i> for your 3400cl switch.

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 in ip from any to any** ACE automatically active after the last operator-created ACE
- ACEs define the ACL for a given client:
 - A given ACE configuration on a RADIUS server includes the identity of the client to which it applies. That is, the ACE includes the client username/password pair or the client device’s MAC address.
 - All ACEs configured on a RADIUS server for the same client are interpreted as belonging to the same ACL. (There is no ACL name or number configured on the RADIUS server.)

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:

VENDOR	HP	11	← ProCurve (HP) Vendor-Specific ID
BEGIN-VENDOR	HP		
ATTRIBUTE	HP-IP-FILTER-RAW	61	← ProCurve (HP) Vendor-Specific Attribute for RADIUS-Based ACLs
END-VENDOR	HP		

Figure 6. Example of Configuring the VSA for RADIUS-Based ACLs in a FreeRADIUS Server

- 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:

<pre>client 10.10.10.125 nastype = other secret = 1234</pre>	<p>Note: The key configured in the switch and the secret configured in the RADIUS server supporting the switch must be identical. Refer to the chapter titled "RADIUS Authentication and Accounting" in the <i>Access Security Guide</i> for your switch.</p>
--	--

Figure 7. Example of Configuring the Switch's Identity Information in a FreeRADIUS Server

- 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 55.) 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 8](#) 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 60.

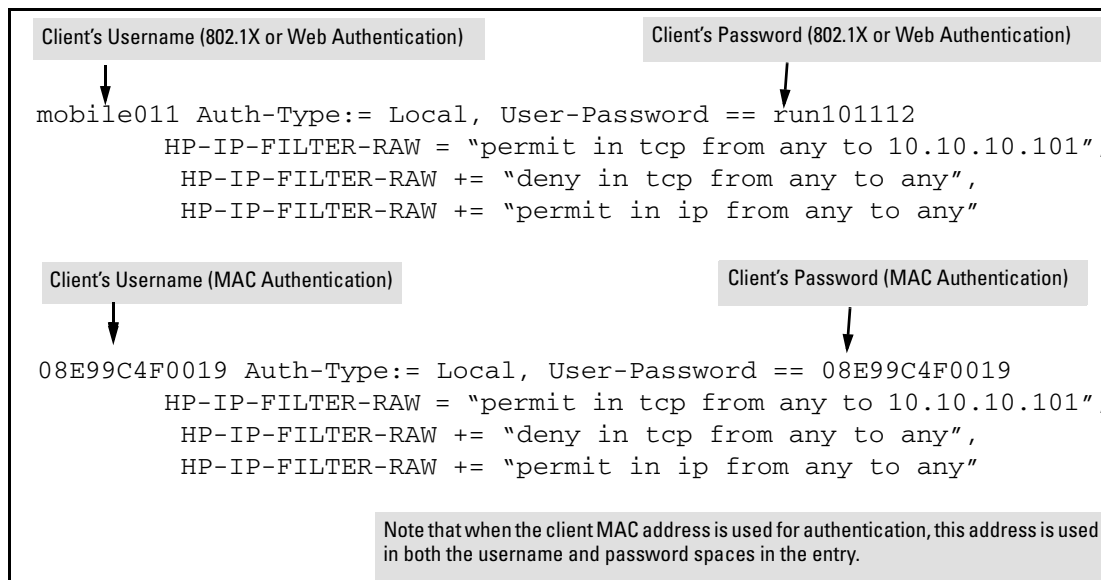


Figure 8. Example of Configuring the FreeRADIUS Server To Support ACLs for the Indicated Clients

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 >] | any > [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	6 = TCP	41 = IPv6
2 = IGMP	17 = UDP	

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. (Refer to the entry describing the maximum number of (optional) internal counters in the table on page 57.) Counter values appear in RADIUS accounting log for client if RADIUS networking accounting is configured on the switch.*

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 >
[auth-port < udp-dest-port > acct-port < udp-dest-port >]

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, including the **auth-port** and **acct-port** options, 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 62) 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 the 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 the following message:*

Port < port-# >, No RADIUS ACLs applied on this port.

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 10 to filter inbound traffic from an authenticated client identified by a MAC address of 00-11-85-C6-54-7D.

```
ProCurve# show access-list radius 10

Radius-configured Port-based ACL for
[Port 10, Client -- 001185C6547D ]
deny in tcp from any to 10.15.240.184 23 cnt
[Packet Hit Counter : 0 ]
deny in tcp from any to 10.15.240.184 80 cnt
[Packet Hit Counter : 0 ]
permit in tcp from any to 10.15.240.184 7
[permit in udp from any to 10.15.240.184 7 ]
deny in tcp from any to 10.15.240.184 161 cnt
Packet Hit Counter : 0
deny in udp from any to 10.15.240.184 161 cnt
Packet Hit Counter : 0
permit in ip from any to any
```

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

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

Lists current counter for the preceding "Deny" ACE.

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

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

Figure 9. 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.

```
ProCurve(config)# show port-access authenticator 10-11

Port Access Authenticator Status

Port-access authenticator activated [No] : No

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

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

Figure 10. 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.
Port <port-#>, No RADIUS ACLs applied on this port.	Appears in response to the CLI show access-list radius <port-#> command when there is not currently a RADIUS ACL assigned to the 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 5 on page 57 .
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 port limit. client <mac-address> port <port-#>.	Notifies that the internal counter (cnt) limit of 32 per port was exceeded on port <port-#>. Refer to Table 5 on page 57 .

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 5, “Limits Affecting RADIUS-Based ACL Applications”](#) on page 57.)
 - The allowed maximum of one RADIUS-assigned ACL 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, the latest client authentication includes a RADIUS-Based ACL assignment exceeding the maximum number of ACEs allowed on the port (30).

- 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 ACL to exceed the per-ACL maximum (32).

SFlow Show Commands

In earlier software releases, the only method for checking whether sFlow is enabled on the switch was via an snmp request. Beginning with software release M.10.02, the 3400cl switches have added the following show sFlow commands that allow you to see sFlow status via the CLI.

Syntax: show sflow agent

Displays sFlow agent information. The agent address is normally the ip address of the first vlan configured.

Syntax: show sflow destination

Displays information about the management station to which the sFlow sampling-polling data is sent.

Syntax: show sflow sampling-polling <port-list/range>

Displays status information about sFlow sampling and polling.

Syntax: show sflow all

Displays sFlow agent, destination, and sampling-polling status information for all the ports on the switch.

Terminology

sFlow — An industry standard sampling technology, defined by RFC 3176, used to continuously monitor traffic flows on all ports providing network-wide visibility into the use of the network.

sFlow agent — A software process that runs as part of the network management software within a device. The agent packages data into datagrams that are forwarded to a central data collector.

sFlow destination — The central data collector that gathers datagrams from sFlow-enabled switch ports on the network. The data collector decodes the packet headers and other information to present detailed Layer 2 to Layer 7 usage statistics.

Viewing SFlow Configuration

The **show sflow agent** command displays read-only switch agent information. The version information shows the sFlow MIB support and software versions; the agent address is typically the ip address of the first vlan configured on the switch.

```
ProCurve# show sflow agent
Version          1.3;HP;M.10.03
Agent Address    10.0.10.228
```

Figure 13. Viewing sFlow Agent Information

The **show sflow destination** command includes information about the management-station's destination address, receiver port, and owner.

```
ProCurve# show sflow destination
sflow           Enabled
Datagrams Sent  221
Destination Address 10.0.10.41
Receiver Port    6343
Owner           admin
Timeout (seconds) 333
Max Datagram Size 1400
Datagram Version Support 5
```

Figure 14. Example of Viewing sFlow Destination Information

Note the following details:

- **Destination Address** remains blank unless it has been configured on the switch via SNMP.
- **Datagrams Sent** shows the number of datagrams sent by the switch agent to the management station since the switch agent was last enabled.
- **Timeout** displays the number of seconds remaining before the switch agent will automatically disable sFlow (this is set by the management station and decrements with time).
- **Max Datagram Size** shows the currently set value (typically a default value, but this can also be set by the management station).

The **show sflow sampling-polling** command displays information about sFlow sampling and polling on the switch. You can specify a list or range of ports for which to view sampling information.

```
ProCurve# show sflow sampling-polling 1-5

sflow destination Enabled

Port  | Sampling                Dropped      | Polling
      | Enabled  Rate          Header  Samples      | Enabled  Interval
-----+-----
1     | Yes      6500000  128    5671234      | Yes      60
2     | No       50        128     0             | Yes     300
3     | Yes      2000     100    24978         | No       30
4     | Yes      200      100    4294967200    | Yes      40
5     | Yes     20000    128     34            | Yes     500
```

Figure 15. Example of Viewing sFlow Sampling and Polling Information

The **show sflow all** command combines the outputs of the preceding three show commands including sFlow status information for all the ports on the switch.

Release M.10.04 Enhancements

Release M.10.04 includes the following enhancements:

- Enhancement (PR_1000330743) - Instrumentation Monitor, which includes Denial of Service (DoS) logging enhancement.
- Enhancement (PR_1000331027) - TCP/UDP port closure enhancement.
- Enhancement (PR_1000330532) - Improved the "show" command display of STP port detail information to assist in monitoring and troubleshooting of the spanning tree protocol.

Instrumentation Monitor

The 3400cl switches have instrumentation to monitor many operating parameters at pre-determined intervals. Beginning with software release M.10.04, this capability can be used to detect anomalies caused by security attacks or other irregular operations on the switch. The following table shows the parameters that can be monitored, and the possible security attacks that may trigger an alert:

Parameter Name	Description
pkts-to-closed-ports	The count of packets per minute sent to closed TCP/UDP ports. An excessive amount of packets could indicate a port scan, in which an attacker is attempting to expose a vulnerability in the switch.
arp-requests	The count of ARP requests processed per minute. A large amount of ARP request packets could indicate an host infected with a virus that is trying to spread itself.

Parameter Name	Description
ip-address-count	The number of destination IP addresses learned in the IP forwarding table. Some attacks fill the IP forwarding table causing legitimate traffic to be dropped.
system-resource-usage (Denial of Service logging)	The percentage of system resources in use. Some Denial-of-Service (DoS) attacks will cause excessive system resource usage, resulting in insufficient resources for legitimate traffic.
login-failures/min	The count of failed CLI login attempts or SNMP management authentication failures. This indicates an attempt has been made to manage the switch with an invalid login or password. Also, it might indicate a network management station has not been configured with the correct SNMP authentication parameters for the switch.
port-auth-failures/min	The count of times a client has been unsuccessful logging into the network
system-delay	The response time, in seconds, of the CPU to new network events such as BPDU packets or packets for other network protocols. Some DoS attacks can cause the CPU to take too long to respond to new network events, which can lead to a breakdown of Spanning Tree or other features. A delay of several seconds indicates a problem.
mac-address-count	The number of MAC addresses learned in the forwarding table. Some attacks fill the forwarding table so that new conversations are flooded to all parts of the network.
mac-moves/min	The average number of MAC address moves from one port to another per minute. This usually indicates a network loop, but can also be caused by DoS attacks.
learn-discards/min	Number of MAC address learn events per minute discarded to help free CPU resources when busy.

Operating Notes

- To generate alerts for monitored events, you must enable the instrumentation monitoring log and/or SNMP trap. The threshold for each monitored parameter is configurable and can be adjusted to minimize false alarms (see [“Configuring Instrumentation Monitor” on page 73](#)).
- When a parameter exceeds its threshold, an alert (event log message and/or SNMP trap) is generated to inform network administrators of this condition. The following example shows an event log message that occurs when the number of MAC addresses learned in the forwarding table exceeds the configured threshold:




Figure 16. Example of Event Log Message generated by Instrumentation Monitor

- Alerts are automatically rate limited to prevent filling the log file with redundant information. The following is an example of alerts that occur when the device is continually subject to the same attack (too many MAC addresses in this instance):

```
W 01/01/90 00:05:00 inst-mon: Limit for MAC addr count (300) is exceeded (321)
W 01/01/90 00:10:00 inst-mon: Limit for MAC addr count (300) is exceeded (323)
W 01/01/90 00:15:00 inst-mon: Limit for MAC addr count (300) is exceeded (322)
W 01/01/90 00:20:00 inst-mon: Limit for MAC addr count (300) is exceeded (324)
W 01/01/90 00:20:00 inst-mon: Ceasing logs for MAC addr count for 15 minutes
```

Figure 17. Example of the rate limiting that occurs when multiple messages are generated

In the preceding example, if a condition is reported 4 times (persists for more than 15 minutes) then alerts cease for 15 minutes. If after 15 minutes the condition still exists, the alerts cease for 30 minutes, then for 1 hour, 2 hours, 4 hours, 8 hours, and after that the persisting condition is reported once a day. Note that ProCurve switches also have the ability to send event log entries to a syslog server.

Known Limitations

As of release M.10.06, the instrumentation monitor runs once every five minutes. The current implementation does not track information such as the port, MAC, and IP address from which an attack is received.

Configuring Instrumentation Monitor

The following commands and parameters are used to configure the operational thresholds that are monitored on the switch. By default, the instrumentation monitor is disabled.

Syntax: [no] instrumentation monitor [parameterName|all] [<low|med|high|limitValue>]

- [log]** : Enables/disables instrumentation monitoring log so that event log messages are generated every time there is an event which exceeds a configured threshold.
(Default threshold setting when instrumentation monitoring is enabled: **enabled**)
- [all]** : Enables/disables all counter types on the switch but does not enable/disable instrumentation monitor logging.
(Default threshold setting when enabled: **see parameter listings below**)
- [arp-requests]** : The number of arp requests that are processed each minute.
(Default threshold setting when enabled: **1000 (med)**)
- [ip-address-count]**: The number of destination IP addresses learned in the IP forwarding table.
(Default threshold setting when enabled: **1000 (med)**)
- [learn-discards]** : The number of MAC address learn events per minute discarded to help free CPU resources when busy.
(Default threshold setting when enabled: **100 (med)**)
- [login-failures]** : The count of failed CLI login attempts or SNMP management authentication failures per hour.
(Default threshold setting when enabled: **10 (med)**)
- [mac-address-count]** : The number of MAC addresses learned in the forwarding table. You must enter a specific value in order to enable this feature.
(Default threshold setting when enabled: **1000 (med)**)
- [mac-moves]** : The average number of MAC address moves per minute from one port to another.
(Default threshold setting when enabled: **100 (med)**)
- [pkts-to-closed-ports]** : The count of packets per minute sent to closed TCP/UDP ports.
(Default threshold setting when enabled: **10 (med)**)
- [port-auth-failures]** : The count of times per minute that a client has been unsuccessful logging into the network.
(Default threshold setting when enabled: **10 (med)**)
- [system-resource-usage]**: The percentage of system resources in use.
(Default threshold setting when enabled: **50 (med)**)
- [system-delay]** : The response time, in seconds, of the CPU to new network events such as BPDU packets or packets for other network protocols.
(Default threshold setting when enabled: **3 seconds (med)**)
- [trap]** : Enables or disables SNMP trap generation.
(Default setting when instrumentation monitoring is enabled: **disabled**)

To enable instrumentation monitor using the default parameters and thresholds, enter the general **instrumentation monitor** command. To adjust specific settings, enter the name of the parameter that you wish to modify, and revise the threshold limits as needed.

Examples

To turn on monitoring and event log messaging with the default medium values:

```
ProCurve(config)# instrumentation monitor
```

To turn off monitoring of the system delay parameter:

```
ProCurve(config)# no instrumentation monitor system-delay
```

To adjust the alert threshold for the MAC address count to the low value:

```
ProCurve(config)# instrumentation monitor mac-address-count low
```

To adjust the alert threshold for the MAC address count to a specific value:

```
ProCurve(config)# instrumentation monitor mac-address-count 767
```

To enable monitoring of learn discards with the default medium threshold value:

```
ProCurve(config)# instrumentation monitor learn-discards
```

To disable monitoring of learn discards:

```
ProCurve(config)# no instrumentation monitor learn-discards
```

To enable or disable SNMP trap generation:

```
ProCurve(config)# [no] instrumentation monitor trap
```

Viewing the Current Instrumentation Monitor Configuration

The **show instrumentation monitor configuration** command displays the configured thresholds for monitored parameters, as shown in [Figure 18](#) on the next page.

An alternate method of determining the current Instrumentation Monitor configuration is to use the **show run** command. However, the show run command output does not display the threshold values for each limit setting.

```
ProCurve# show instrumentation monitor configuration

PARAMETER                                LIMIT
-----
mac-address-count                        1000 (med)
ip-address-count                          1000 (med)
system-resource-usage                     50 (med)
system-delay                              5 (high)
mac-moves/min                             100 (med)
learn-discards/min                        100 (med)
ip-port-scans/min                         10 (med)
arp-requests/min                          100 (low)
login-failures/min                        10 (med)
port-auth-failures/min                    10 (med)

SNMP trap generation for alerts: enabled
Instrumentation monitoring log : enabled
```

Figure 18. Viewing the Instrumentation Monitor Configuration

TCP/UDP Port Closure

In earlier software releases, certain UDP ports were always open. Beginning with software release M.10.04, all TCP/UDP ports on the 3400cl switches will remain closed until the associated services are enabled on the switch.

The following ports and services are affected by this change:

Port	Service
69	TFTP
161	SNMP
520	RIP
1507	Stacking (SNMP)

To open any of these ports, the respective services must first be enabled on the switch. For information on how to enable/disable these services, refer to the following command listings . For details on each service, refer to the latest version of the switch's software documentation available on the ProCurve Networking Web site.

Enabling/Disabling TFTP

The TFTP server and client can be enabled and/or disabled independently.

Syntax: [no] tftp < client | server >

Enables or disables the TFTP client.

client: *Enables or disables the TFTP client.*

(Default: disabled)

server: *Enables or disables the TFTP server.*

(Default: disabled)

Note: Both the **tftp** command (with no arguments) and the **tftp client** command can be used to enable or disable the tftp client.

Enabling/Disabling SNMP

To enable/disable SNMP, use the following commands.

Syntax: [no] snmp-server enable

Enables or disables SNMP v1/v2.

(Default: disabled)

Syntax: [no] snmpv3 enable

Enables or disables SNMP v3.

(Default: disabled)

Notes

- The SNMP port (161) will be opened if either SNMP v1/2 or SNMP v3 are enabled, or remain closed if both are disabled.
- The **snmp-server enable** command takes precedence over the **snmp-server enable traps** command that is used to enable or disable authentication traps to be sent when a management station attempts an unauthorized access.
- If SNMP is disabled, both the SNMP port (161) and the stacking port (1507) will remain closed.

Enabling/Disabling RIP

To enable/disable RIP, use the following command.

Syntax: [no] router rip

Enables, disables, or configures Routing Internet Protocol (RIP) on the switch.

(Default: disabled)

Note

The **router rip** command exists in previous software versions. In this implementation, however, RIP must be enabled in order to open the port on the switch.

Enabling/Disabling Stacking

To enable/disable stacking, use the following command.

Syntax: [no] stack

Enables stacking (SNMP) on the switch. (Default: disabled)

Note

The **stack** command exists in previous software versions. In this implementation, however, both stacking and SNMP must be enabled to open the port on the switch. If either feature is disabled, the port will remain closed.

Spanning Tree Show Commands

The **show spanning-tree detail** command previously displayed 802.1D (STP) and 802.1w (RSTP) status and counters for all ports on the switch. Beginning with software release M.10.04, this command provides 802.1s (MSTP) multi-instance spanning tree details and displays additional parameters to enhance spanning-tree reporting via the CLI.

The following shows RSTP sample output from the enhanced command.

```
ProCurve# show spanning-tree detail

Status and Counters - RSTP Port(s) Detailed Information

Port                : 1
Status              : Up
Role                : Root
State               : Forwarding
Priority             : 128
Path Cost           : 200000
Root Path Cost      : 10
Root Bridge ID      : 1:0001e7-215e00
Designated Bridge ID : 32768:0001e7-3d0080
Designated Port ID  : 128:75
AdminEdgePort       : Yes
OperEdgePort        : No
AdminPointToPointMAC : Force-True
OperPointToPointMAC  : Yes
Aged BPDUs Count    : 0
Loop-back BPDUs Count : 0
TC Detected          : 1
TC Flag Transmitted  : 0          TC ACK Flag Transmitted : 0
TC Flag Received     : 0          TC ACK Flag Received   : 47

RSTP      RSTP      CFG      CFG      TCN      TCN
BPDUs Tx  BPDUs Rx  BPDUs Tx  BPDUs Rx  BPDUs Tx  BPDUs Rx
-----
3          0          0          256654    47         0
```

Figure 19. Example of Show Spanning-Tree Detail

Operating Notes

- TC refers to a Topology Change detected on the given port. Note the following details:
 - **TC Detected** counter shows when a port identifies a topology change (increments when the particular non-Edge port goes into forwarding). For RSTP and MSTP, this would be due to the switch's link going to forwarding.
 - **TC Flag Transmitted** counter shows the number of TC notifications sent out of the port. This refers to propagating a topology change that occurred on another port (that is, a TC Detected increment) or to propagating a topology change received on another port (that is, TC Flag Received).

- **TC Flag Received** counter shows the number of TC notifications (RSTP or MSTP style BPDU with the TC flag set) received on the port.
 - **TC ACK Flag Transmitted** is an 802.1D mode counter. It will only increment when the port is operating in 802.1D mode and an 802.1D style PDU is sent out of the port.
 - **TC ACK Flag Received** is an 802.1D mode counter. It will only increment when the port is operating in 802.1D mode and an 802.1D style PDU is received on the port.
- With STP and RSTP activated:
- The **show spanning tree detail** command shows all active RSTP port by port.
 - The **show spanning-tree <port-list> detail** command shows the specified port-list RSTP port by port detail.
- With MSTP activated:
- The **show spanning tree detail** command shows all active MSTP port by port. This command only gives information concerning the common spanning tree (CST) ports. To view counters pertaining to a specific spanning-tree instance, you must use the **show spanning-tree instance <inst> detail** command. The **show spanning-tree <port-list> detail** command shows the specified port-list MSTP port by port detail.
 - The **show spanning-tree instance <inst> detail** command shows all ports active for a specific instance of MSTP.
 - The **show spanning-tree <port-list> instance <inst> detail** shows the specified port-list for the specified instance of MSTP.
 - **TC ACK Flag Transmitted** and **TC ACK Flag Received** are part of the CST counters displayed by the **show spanning tree detail** command. **TC Detected**, **TC Flag Transmitted**, and **TC Flag Received** are included only with the **instance** parameter due to the nature of MSTP.
-

Release M.10.05 Enhancements

Release M.10.05 includes the following enhancement:

- Ping functionality now in conformance with RFC 2925 specification.
-

Release M.10.06 Enhancements

Release M.10.06 includes the following enhancement:

- Enhancement (PR_1000330704) - Added RADIUS Command Authorization and Accounting for the Command Line Interface on 3400cl switch. Please refer to Chapter 6, RADIUS Authentication and Accounting in the *Access Security Guide for the ProCurve Series 6400cl/5300xl/4200vl/3400cl Switches* (October 2005) for additional information.
-

Release M.10.07 Enhancements

Release M.10.07 includes the following enhancement:

- Added support for PIM Dense Mode. For details, refer to Chapter 5, “PIM-DM (Dense Mode) on the 5300xl Switches” in the *Advanced Traffic Management Guide for the ProCurve Series 6400cl/5300xl/4200vl/3400cl Switches*.

Release M.10.08 Enhancements

Software fixes only, no new enhancements.

Release M.10.09 Enhancements

Release M.10.09 includes the following enhancement:

- Added support for Unidirectional Link Detection (UDLD). See [“Uni-Directional Link Detection \(UDLD\)” on page 80](#) for details.

Uni-Directional Link Detection (UDLD)

Uni-directional Link Detection (UDLD) monitors a link between two ProCurve switches and blocks the ports on both ends of the link if the link fails at any point between the two devices. This feature is particularly useful for detecting failures in fiber links and trunks.

In the example shown in [Figure 20](#), each ProCurve switch load balances traffic across two ports in a trunk group. Without the UDLD feature, a link failure on a link that is not directly attached to one of the ProCurve switches remains undetected. As a result, each switch continues to send traffic on the ports connected to the failed link. When UDLD is enabled on the trunk ports on each ProCurve switch, the switches detect the failed link, block the ports connected to the failed link, and use the remaining ports in the trunk group to forward the traffic.

Scenario 1 (No UDLD): Without UDLD, the switch ports remain enabled despite the link failure. Traffic continues to be load-balanced to the ports connected to the failed link.

Scenario 2 (UDLD-enabled): When UDLD is enabled, the feature blocks the ports connected to the failed link.

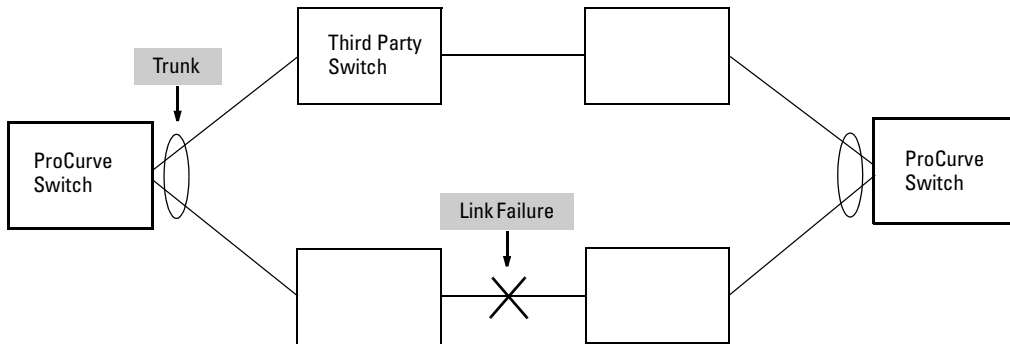


Figure 20. UDLD Example

Similarly, UDLD is effective for monitoring fiber optic links that use two uni-direction fibers to transmit and receive packets. Without UDLD, if a fiber breaks in one direction, a fiber port may assume the link is still good (because the other direction is operating normally) and continue to send traffic on the connected ports. UDLD-enabled ports, however, prevent traffic from being sent across a bad link by blocking the ports in the event that either the individual transmitter or receiver for that connection fails.

Ports enabled for UDLD exchange health-check packets once every five seconds (the link-keepalive interval). If a port does not receive a health-check packet from the port at the other end of the link within the keepalive interval, the port waits for four more intervals. If the port still does not receive a health-check packet after waiting for five intervals, the port concludes that the link has failed and blocks the UDLD-enabled port.

When a port is blocked by UDLD, the event is recorded in the switch log or via an SNMP trap (if configured); and other port blocking protocols, like spanning tree or meshing, will not use the bad link to load balance packets. The port remains blocked until the link is unplugged, disabled, or fixed. The port can also be unblocked by disabling UDLD on the port.

Configuration Considerations

- UDLD is configured on a per-port basis and must be enabled at both ends of the link. See the note below for a list of ProCurve switches that support UDLD.
- To configure UDLD on a trunk group, you must configure the feature on each port of the group individually. Configuring UDLD on a trunk group's primary port enables the feature on that port only.
- Dynamic trunking is not supported. If you want to configure a trunk group that contains ports on which UDLD is enabled, you must remove the UDLD configuration from the ports. After you create the trunk group, you can re-add the UDLD configuration.

Note

UDLD interoperates with the following ProCurve switch series: 2600, 2800, 3400, 3500, 4200, 5300, 5400, 6200, 6400, and 9300. Consult the release notes and current manuals for required software versions.

Configuring UDLD

The following commands allow you to configure UDLD via the CLI.

Syntax: [no] interface <port-list> link-keepalive

Enables UDLD on a port or range of ports.

*To disable the feature, enter the **no** form of the command.*

Default: UDLD disabled

Syntax: link-keepalive interval <interval>

Determines the time interval to send UDLD control packets. The <interval> parameter specifies how often the ports send a UDLD packet. You can specify from 10 – 100, in 100 ms increments, where 10 is 1 second, 11 is 1.1 seconds, and so on.

Default: 50 (5 seconds)

Syntax: link-keepalive retries <num>

Determines the maximum number of retries to send UDLD control packets. The <num> parameter specifies the maximum number of times the port will try the health check. You can specify a value from 3 – 10.

Default: 5

Syntax: [no] interface <port-list> link-keepalive vlan <vid>

Assigns a VLAN ID to a UDLD-enabled port for sending of tagged UDLD control packets. Under default settings, untagged UDLD packets can still be transmitted and received on tagged only ports—however, a warning message will be logged.

*The **no** form of the command disables UDLD on the specified port(s).*

Default: UDLD packets are untagged; tagged only ports will transmit and receive untagged UDLD control packets

Enabling UDLD. UDLD is enabled on a per port basis. For example, to enable UDLD on port a1, enter:

```
ProCurve(config)#interface a1 link-keepalive
```

To enable the feature on a trunk group, enter the appropriate port range. For example:

```
ProCurve(config)#interface a1-a4 link-keepalive
```

Note

When at least one port is UDLD-enabled, the switch will forward out UDLD packets that arrive on non-UDLD-configured ports out of all other non-UDLD-configured ports in the same vlan. That is, UDLD control packets will “pass through” a port that is not configured for UDLD. However, UDLD packets will be dropped on any blocked ports that are not configured for UDLD.

Changing the Keepalive Interval. By default, ports enabled for UDLD send a link health-check packet once every 5 seconds. You can change the interval to a value from 10 – 100 deciseconds, where 10 is 1 second, 11 is 1.1 seconds, and so on. For example, to change the packet interval to seven seconds, enter the following command at the global configuration level:

```
ProCurve(config)# link-keepalive interval 70
```

Changing the Keepalive Retries. By default, a port waits five seconds to receive a health-check reply packet from the port at the other end of the link. If the port does not receive a reply, the port tries four more times by sending up to four more health-check packets. If the port still does not receive a reply after the maximum number of retries, the port goes down.

You can change the maximum number of keepalive attempts to a value from 3 – 10. For example, to change the maximum number of attempts to 4, enter the following command at the global configuration level:

```
ProCurve(config)# link-keepalive retries 4
```

Configuring UDLD for Tagged Ports. The default implementation of UDLD sends the UDLD control packets untagged, even across tagged ports. If an untagged UDLD packet is received by a non-ProCurve switch, that switch may reject the packet. To avoid such an occurrence, you can configure ports to send out UDLD control packets that are tagged with a specified VLAN.

To enable ports to receive and send UDLD control packets tagged with a specific VLAN ID, enter a command such as the following at the interface configuration level:

```
ProCurve(config)#interface 1 link-keepalive vlan 22
```

Notes

- You must configure the same VLANs that will be used for UDLD on all devices across the network; otherwise, the UDLD link cannot be maintained.
 - If a VLAN ID is not specified, then UDLD control packets are sent out of the port as untagged packets.
 - To re-assign a VLAN ID, re-enter the command with the new VLAN ID number. The new command will overwrite the previous command setting.
 - When configuring UDLD for tagged ports, you may receive a warning message if there are any inconsistencies with the port's VLAN configuration (see page 87 for potential problems).
-

Viewing UDLD Information

The following show commands allow you to display UDLD configuration and status via the CLI.

Syntax: show link-keepalive

Displays all the ports that are enabled for link-keepalive.

Syntax: show link-keepalive statistics

Displays detailed statistics for the UDLD-enabled ports on the switch.

Syntax: clear link-keepalive statistics

Clears UDLD statistics. This command clears the packets sent, packets received, and transitions counters in the show link-keepalive statistics display.

Displaying Summary UDLD Information. To display summary information on all UDLD-enabled ports, enter the **show link-keepalive** command. For example:

```
ProCurve(config)# show link-keepalive

Total link-keepalive enabled ports: 4
Keepalive Retries: 3           Keepalive Interval: 1 sec

Port Enabled Physical  Keepalive  Adjacent  UDLD
      Status Status      Status      Switch    VLAN
-----
1  Yes   up      up        00d9d-f9b700  200
2  Yes   up      up        01560-7b1600
3  Yes   down   off-line
4  Yes   up      failure
5  No    down   off-line
```

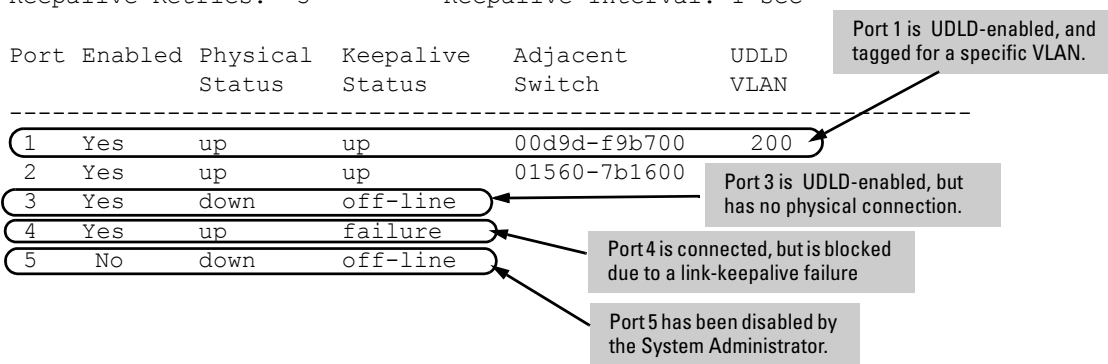


Figure 21. Example of UDLD Information displayed using Show Link-Keepalive Command

Displaying Detailed UDLD Status Information. To display detailed UDLD information for specific ports, enter the **show link-keepalive statistics** command. For example:

```
ProCurve(config)# show link-keepalive statistics
```

Port:	1	Neighbor MAC Addr:	0000a1-b1c1d1
Current State:	up	Neighbor Port:	5
Uddl Packets Sent:	1000	State Transitions:	2
Uddl Packets Received:	1000	Link-vlan:	1
Port Blocking:	no		
Port:	2	Neighbor MAC Addr:	000102-030405
Current State:	up	Neighbor Port:	6
Uddl Packets Sent:	500	State Transitions:	3
Uddl Packets Received:	450	Link-vlan:	200
Port Blocking:	no		
Port:	3	Neighbor MAC Addr:	n/a
Current State:	off line	Neighbor Port:	n/a
Uddl Packets Sent:	0	State Transitions:	0
Uddl Packets Received:	0	Link-vlan:	1
Port Blocking:	no		
Port:	4	Neighbor MAC Addr:	n/a
Current State:	failure	Neighbor Port:	n/a
Uddl Packets Sent:	128	State Transitions:	8
Uddl Packets Received:	50	Link-vlan:	1
Port Blocking:	yes		

Ports 1 and 2 are UDLD-enabled and show the number of health check packets sent and received on each port.

Port 4 is shown as blocked due to a link-keepalive failure

Figure 22. Example of Detailed UDLD Information displayed using Show Link-Keepalive Statistics Command

Clearing UDLD Statistics. To clear UDLD statistics, enter the following command:

```
ProCurve# clear link-keepalive statistics
```

This command clears the Packets sent, Packets received, and Transitions counters in the **show link keepalive statistics** display (see [Figure 22](#) for an example).

Configuration Warnings and Event Log Messages

Warning Messages. The following table shows the warning messages that may be issued and their possible causes, when UDLD is configured for tagged ports.

Table 6. Warning Messages caused by configuring UDLD for Tagged Ports

CLI Command Example	Warning Message	Possible Problem
link-keepalive 6	Possible configuration problem detected on port 6. UDLD VLAN configuration does not match the port's VLAN configuration.	You have attempted to enable UDLD on a port that is a tagged only port, but did not specify a configuration for tagged UDLD control packets. In this example, the switch will send and receive the UDLD control packets untagged despite issuing this warning.
link-keepalive 7 vlan 4	Possible configuration problem detected on port 7. UDLD VLAN configuration does not match the port's VLAN configuration.	You have attempted to configure tagged UDLD packets on a port that does not belong to the specified VLAN. In this example, if port 7 belongs to VLAN 1 and 22, but the user tries to configure UDLD on port 7 to send tagged packets in VLAN 4, the configuration will be accepted. The UDLD control packets will be sent tagged in VLAN 4, which may result in the port being blocked by UDLD if the user does not configure VLAN 4 on this port.
no vlan 22 tagged 20	Possible configuration problem detected on port 18. UDLD VLAN configuration does not match the port's VLAN configuration.	You have attempted to remove a VLAN on port that is configured for tagged UDLD packets on that VLAN. In this example, if port 18, 19, and 20 are transmitting and receiving tagged UDLD packets for Vlan 22, but the user tries to remove Vlan 22 on port 20, the configuration will be accepted. In this case, the UDLD packets will still be sent on Vlan 20, which may result in the port being blocked by UDLD if the users do not change the UDLD configuration on this port.

Note: If you are configuring the switch via SNMP with the same problematic VLAN configuration choices, the above warning messages will also be logged in the switch's event log.

Event Log Messages. The following table shows the event log messages that may be generated once UDLD has been enabled on a port.

Table 7. UDLD Event Log Messages

Message	Event
I 01/01/06 04:25:05 ports: port 4 is deactivated due to link failure.	A UDLD-enabled port has been blocked due to part of the link having failed.
I 01/01/06 06:00:43 ports: port 4 is up, link status is good.	A failed link has been repaired and the UDLD-enabled port is no longer blocked.

Release M.10.10 Enhancements

Release M.10.10 includes the following enhancement:

Spanning Tree Per-Port BPDU Filtering

The STP BPDU filter feature allows control of spanning-tree participation on a per-port basis. It can be used to exclude specific ports from becoming part of spanning tree operations. A port with the BPDU filter enabled will ignore incoming BPDU packets and stay locked in the spanning-tree forwarding state. All other ports will maintain their role.

Here are some sample scenarios in which this feature may be used:

- To have STP operations running on selected ports of the switch rather than every port of the switch at a time.
- To prevent the spread of errant BPDU frames.
- To eliminate the need for a topology change when a port's link status changes. For example, ports that connect to servers and workstations can be configured to remain outside of standard spanning-tree operations.
- To protect the network from denial of service attacks with spoofing spanning-tree BPDUs by dropping incoming BPDU frames.

Note

BPDU protection imposes a more secure mechanism that implements port shut down and a detection alert when an errant BPDU frame is received ([see page 91](#) for details). BPDU protection will take precedence over BPDU filtering if both features have been enabled on the same port.

Configuring STP BPDU Filters

The following commands allow you to configure BPDU filters via the CLI.

Syntax: [no] spanning-tree <port-list | all> bpdu-filter

Enables/disables the BPDU filter feature on the specified port(s).

For example, to configure BPDU filtering on port a9, enter:

```
ProCurve(config)# spanning-tree a9 bpdu-filter
```

Caution

Ports configured with the BPDU filter mode remain active (learning and forward frames); however, spanning-tree cannot receive or transmit BPDUs on the port. The port remains in a forwarding state, permitting all broadcast traffic. This can create a network storm if there are any loops (that is, trunks or redundant links) using these ports. If you suddenly have a high load, disconnect the link and remove ("no") the bpdud-filter.

Viewing Status of BPDU Filtering

The **show spanning-tree <port-list> detail** command has been extended to show per-port BPDU filter mode as shown below.

```

ProCurve# show spanning-tree a9 detail

Status and Counters - CST Port(s) Detailed Information

Port                : A1
Status              : Up
BPDU Filtering      : Yes
Errant BPUDUs received : 65
MST Region Boundary : Yes
External Path Cost  : 200000
External Root Path Cost : 420021
Administrative Hello Time : Use Global
Operational Hello Time   : 2
AdminEdgePort          : No
OperEdgePort           : No
AdminPointToPointMAC    : Force-True
OperPointToPointMAC     : Yes
Aged BPDUs Count        : 0
Loop-back BPDUs Count   : 0
TC ACK Flag Transmitted : 0
TC ACK Flag Received    : 0

MST          MST          CFG          CFG          TCN          TCN
BPDU Tx     BPDU Rx     BPDU Tx     BPDU Rx     BPDU Tx     BPDU Rx
-----
8           28          0           0           0           0
  
```

Figure 23. Example of BPDU Filter Fields in Show Spanning Tree Detail Command

The **show spanning-tree** command has also been extended to display BPDU filtered ports.

```
ProCurve# show spanning-tree

Multiple Spanning Tree (MST) Information

STP Enabled    : Yes
Force Version  : MSTP-operation
IST Mapped VLANs : 1-7
...

Protected Ports :
Filtered Ports  : A6-A7
....
```

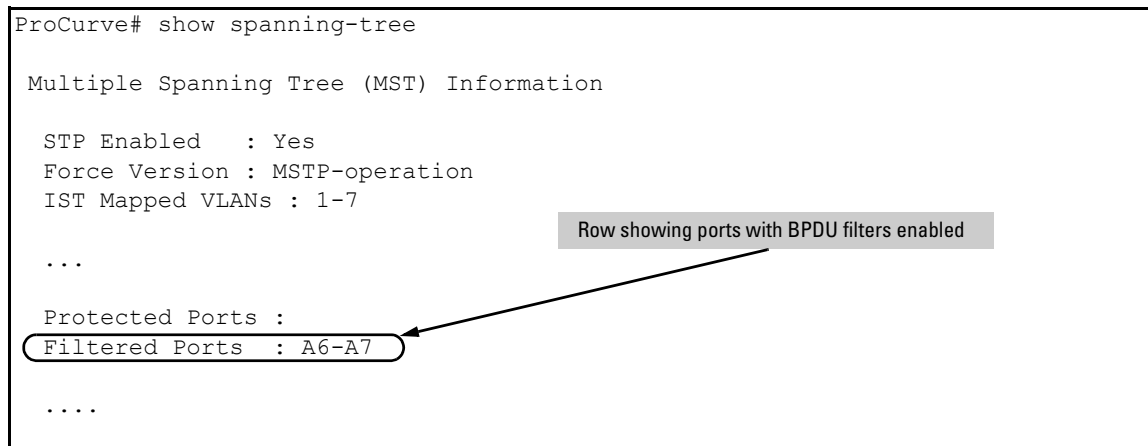


Figure 24. Example of BPDU Filtered Ports Field in Show Spanning Tree Command

Viewing Configuration of BPDU Filtering

The BPDU filter mode adds an entry to the spanning tree category within the configuration file.

```
ProCurve(config)# show configuration
. . .
spanning-tree
spanning-tree A7 bpdu-filter
spanning-tree C9 bpdu-filter
spanning-tree Trk2 priority 4
. . .
```

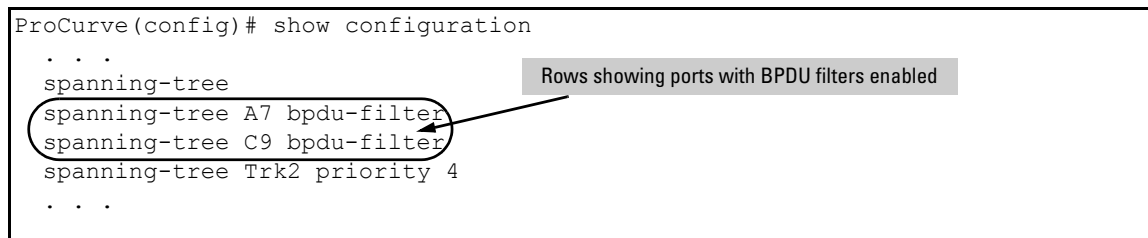


Figure 25. Example of BPDU Filters in the Show Configuration Command

The **spanning-tree show < port> configuration** command displays the BPDU's filter state.

```
ProCurve(config)# show spanning-tree a8 config

...

Port Type      | Cost      Priority Edge Point-to-Point MCheck Filter
-----+-----
A8 100/1000T | Auto      128     Yes  Force-True   Yes   No
```

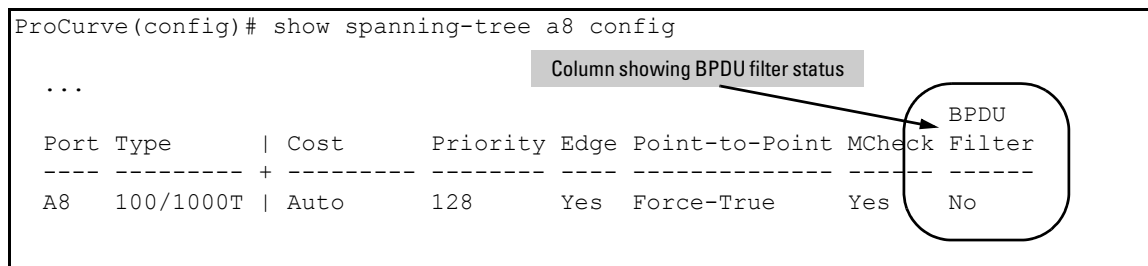


Figure 26. Example of BPDU Filter Status in Show Spanning Tree Configuration Command

Releases M.10.11 through M.10.12 Enhancements

Software fixes only, no new enhancements.

Release M.10.13 Enhancements

Release M.10.13 includes the following enhancement:

- Enhancement (PR_1000354065) - Added DHCP protection feature. No additional documentation is available at this time
-

Releases M.10.14 through M.10.16 Enhancements

Software fixes only, no new enhancements.

Release M.10.17 Enhancements

Release M.10.17 includes the following enhancement:

- **RSTP/MSTP BPDU Protection.** When this feature is enabled on a port, the switch will disable (drop the link) of a port that receives a spanning tree BPDU, log a message, and optionally, send an SNMP trap.

Spanning Tree BPDU Protection

The BPDU protection feature is a security enhancement to Spanning Tree Protocol (STP) operation. It can be used to protect the active STP topology by delimiting its legal boundaries, thereby preventing spoofed BPDU packets from entering the STP domain. In a typical implementation, BPDU protection would be applied to edge ports connected to end user devices that do not run STP. If STP BPDU packets are received on a protected port, the feature will disable that port and alert the network manager via an SNMP trap as shown in [Figure 27](#).

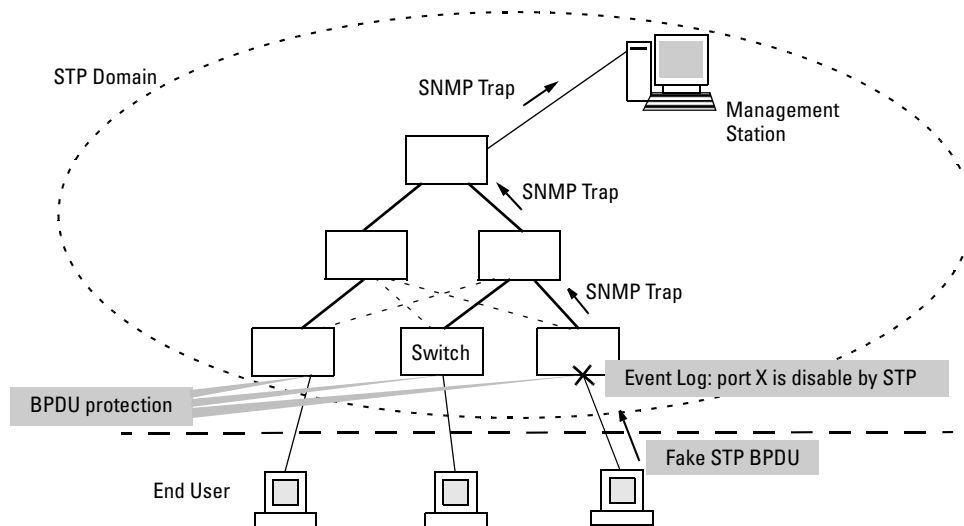


Figure 27. Example of BPDU Protection Enabled at the Network Edge

Terminology

BPDU — Acronym for bridge protocol data unit. BPDUs are data messages that are exchanged between the switches within an extended LAN that use a spanning tree protocol topology. BPDU packets contain information on ports, addresses, priorities and costs and ensure that the data ends up where it was intended to go. BPDU messages are exchanged across bridges to detect loops in a network topology. The loops are then removed by placing redundant switch ports in a backup, or blocked, state.

BPDU Filtering — Spanning-tree configuration mode that prevents the switch from receiving and transmitting BPDU frames on a specific port.

BPDU Protection — Spanning-tree configuration mode which disables a port where BPDU frames are received.

MSTP — Multiple Spanning Tree Protocol, defined in IEEE 802.1s. Each MSTI (multiple spanning tree instance) on a physical port provides loop free connectivity for the group of VLANs associated with that instance. This means that traffic transported on different VLANs can be distributed for load-balancing among links between switches.

RSTP — Rapid Spanning Tree Protocol, defined in IEEE 802.1w and ratified in IEEE 802.1D-2004.

Spanning-tree — Generic term to refer to the many spanning-tree flavors: now deprecated STP, RSTP and VLAN-aware MSTP.

STP — Spanning Tree Protocol, part of the original IEEE 802.1D specification. The 2004 edition completely deprecates STP. Both RSTP and MSTP have fallback modes to handle STP.

SNMP — Simple Network Management Protocol, used to remotely manage network devices.

Note

The switches covered in these Release Notes, use the IEEE 802.1s Multiple Spanning Tree Protocol (MSTP) standard. Under standard settings, your MSTP-configured switch interoperates effectively with both STP (IEEE 802.1D) and RSTP (IEEE 802.1w) spanning-tree devices. For more information, refer to the chapter entitled *Multiple Instance Spanning-Tree Operation* in the *Advanced Traffic Management Guide* for your switch.

Configuring STP BPDU Protection

The following commands allow you to configure BPDU protection via the CLI.

Syntax: [no] spanning-tree <port-list> bpdu protection

Enables/disables the BPDU protection feature on a port

Syntax: [no] spanning-tree trap errant bpdu

Enables/disables the sending of errant BPDU traps.

For example, to configure BPDU protection on ports 1 to 10, enter:

```
ProCurve(config)# spanning-tree 1-10 bpdu protection
```

When BPDU protection is enabled, the following steps are set in process:

1. When an STP BPDU packet is received, STP treats it as an unauthorized transmission attempt and shuts down the port that the BPDU came in on.
2. An event message is logged and an SNMP notification trap is generated.
3. The port remains disabled until re-enabled manually by a network administrator.

Caution

This command should only be used to guard edge ports that are not expected to participate in STP operations. Once BPDU protection is enabled, it will disable the port as soon as any BPDU packet is received on that interface.

Viewing BPDU Protection Status

The **show spanning-tree** command has additional information on BPDU protection as shown below.

```
ProCurve# show spanning-tree 1-10

Multiple Spanning Tree (MST) Information

STP Enabled      : Yes
Force Version    : MSTP-operation
IST Mapped VLANs : 1-7

...

Protected Ports : 3-7,9
Filtered Ports  : 10
```

Port	Type	Cost	Priority	State	Designated Bridge	Hello Time	PtP	Edge
1	100/1000T	200000	128	Forwarding	000883-024500	2	Yes	No
2	100/1000T	200000	128	Forwarding	000883-122740	2	Yes	No
3	100/1000T	200000	128	BpduError		2	Yes	Yes
4	100/1000T	Auto	128	Disabled				
5	100/1000T	200000	128	Forwarding		2	Yes	Yes
6	100/1000T	200000	128	Forwarding		2	Yes	Yes
7	100/1000T	200000	128	Forwarding		2	Yes	Yes
8	100/1000T	Auto	128	Disabled				
9	100/1000T	Auto	128	Disabled				
10	100/1000T	200000	128	Forwarding		2	Yes	Yes

Annotations in the image:

- A grey box labeled "Ports with BPDU protection enabled" has an arrow pointing to the "Protected Ports : 3-7,9" line.
- A grey box labeled "Errant BPDU detected on this port" has an arrow pointing to the "BpduError" state in the table row for port 3.
- A white box with a black border labeled "Protected Ports : 3-7,9" has an arrow pointing to the "Protected Ports" line.

Example of BPDU Protection Additions to Show Spanning Tree Command

Release M.10.21 Enhancements

Software fixes only, no new enhancements.

Release M.10.22 Enhancements

Release M.10.22 includes the following enhancement:

- **Enhancement (PR_1000376406)** — Loop Protection feature additions, including packet authentication, loop detected trap, and receiver port configuration.

Configuring Loop Protection

You can use BPDU protection for systems that have spanning tree enabled (See [“Spanning Tree BPDU Protection” on page 91](#)), however, the BPDU protection feature cannot detect the formation of loops when an unmanaged device on the network drops spanning tree packets. To protect against the formation of loops in these cases, you can enable the Loop Protection feature, which provides protection by transmitting loop protocol packets out ports on which loop protection has been enabled. When the switch sends out a loop protocol packet and then receives the same packet on a port that has **send-disable** configured, it shuts down the port from which the packet was sent.

You can configure the **disable-timer** parameter for the amount of time you want the port to remain disabled (0 to 604800 seconds). If you configure a value of zero, the port will not be re-enabled.

To enable loop protection, enter this command:

```
ProCurve(config)# loop-protect <port-list>
```

Syntax: [no] loop-protect <port-list> [receiver-action <send-disable | no-disable> |]
[transmit-interval <1-10>] | [disable-timer <0-604800>] |
[trap <loop-detected>]

Allows you to configure per-port loop protection on the switch.

[receiver-action <send-disable | no-disable>]

Sets the action to be taken when a loop is detected on the port. The port that received the loop protection packet determines what action is taken. If send-disable is configured, the port that transmitted the packet is disabled. If no-disable is configured, the port is not disabled.

Default: send-disable

[trap <loop-detected>]

Allows you to configure loop protection traps. The "loop-detected" trap indicates that a loop was detected on a port.

[disable-timer <0-604800>]

How long (in seconds) a port is disabled when a loop has been detected. A value of zero disables the auto re-enable functionality.

Default: Timer is disabled

[transmit-interval <1-10>]

Allows you to configure the time in seconds between the transmission of loop protection packets.

Default: 5 seconds

To display information about ports with loop protection, enter this command.

Syntax: show loop-protect <port-list>

Displays the loop protection status. If no ports are specified, the information is displayed only for the ports that have loop protection enabled.

```
ProCurve(config)# show loop-protect 1-4
```

```
Status and Counters - Loop Protection Information
```

```
Transmit Interval (sec) : 5  
Port Disable Timer (sec) : 5  
Loop Detected Trap      : Enabled
```

Port	Loop Protection	Loop Detected	Loop Count	Time Since Last Loop	Rx Action	Port Status
1	Yes	No	0		send-disable	Up
2	Yes	No	0		send-disable	Up
3	Yes	No	0		send-disable	Up
4	Yes	No	0		send-disable	Up

Figure 28. Example of Show Loop Protect Display

Release M.10.23 Enhancements

Release M.10.23 includes the following enhancement:

- **Enhancement (PR_1000379804)** — Historical information about MAC addresses that have been moved has been added to the "show tech" command output.

Release M.10.24 Enhancements

Release M.10.24 includes the following enhancement:

- **Enhancement (PR_1000335860)** — This enhancement provides a configuration option for the source IP address field of SNMP response and generated trap PDUs.

Release M.10.25 Enhancements

Release M.10.25 includes the following enhancement:

- **Enhancement (PR_1000385565)** — (CLI) The port security MAC address limit per port has been increased from 8 to 32 when learn mode is 'static' or 'configured'. However, the global limit of static/configured MAC addresses per ProCurve Series 3400 switch is 400.

Release M.10.26 Enhancements

Release M.10.26 includes the following enhancement:

- **Enhancement (PR_1000381681)** — This enhancement added eavesdrop protection - the ability to filter unknown Destination IP Address (DA) traffic. For more information, refer to "Eavesdrop Protection" in the chapter "Configuring and Monitoring Port Security" of the *Access Security Guide* for this product.

Release M.10.27 Enhancements

Release M.10.27 includes the following enhancement:

- **Enhancement (PR_1000374085)** — This enhancement expands the use of the Controlled Directions parameter to also support MAC/Web authentication.

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

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

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

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

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

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

*For information on how to configure the prerequisites for using the **aaa port-access controlled-directions in** command, see the chapter titled “Spanning-Tree Operation” in the Advanced Traffic Management Guide for your switch.*

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

Notes:

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

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

- *Using the **aaa port-access controlled-directions in** command, you can enable the transmission of Wake-on-LAN traffic on unauthenticated outbound ports that are configured for any of the following port-based security features:*

- *802.1X authentication*
- *MAC authentication*
- *Web authentication*

*Because a port can be configured for more than one type of authentication to protect the switch from unauthorized access, the last setting you configure with the **aaa port-access controlled-directions** command is applied to all authentication methods configured on the switch.*

For information about how to configure and use 802.1X authentication, refer to the chapter titled “Configuring Port-Based and Client-Based Access Control (802.1X)” in the Access Security Guide for your switch model.

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

Release M.10.28 Enhancements

Software fixes only, no new enhancements.

Release M.10.29 Enhancements

Release M.10.29 includes the following enhancement:

- **Enhancement (PR_1000376626)** — Enhance CLI "qos dscp-map he" help and "show dscp-map" text to warn the user that inbound classification based on DSCP codepoints only occurs if "qos type-of-service diff-services" is also configured.
-

Release M.10.30 Enhancements

Software fixes only, no new enhancements.

Release M.10.31 Enhancements

Release M.10.31 includes the following enhancement:

- **Enhancement (PR_1000372989)** — This enhancement enables the user to set the operator/manager username/password via SNMP.

Password Command

The `password` command in the CLI is enhanced to support the following syntax:

Syntax:

```
[no] password <manager | operator | port-access> [user-name <name>] <hash-type>  
<password>
```

Where:

- **manager** configures access to the switch with manager-level privileges.
 - **operator** configures access to the switch with operator-level privileges.
 - **port-access** configures access to the switch through 802.1X authentication with operator-level privileges.
 - **user-name** <name> is the (optional) text string of the user name associated with the password.
-

- The `<hash-type>` parameter specifies the type of algorithm (if any) used to hash the password. Valid values are **plaintext** or **sha-1**.
- The `<password>` parameter is the clear ASCII text string or SHA-1 hash of the password.

You can enter a manager, operator, or 802.1X port-access password in clear ASCII text or hashed format. However, manager and operator passwords are displayed and saved in a configuration file only in hashed format; port-access passwords are displayed and saved only as plain ASCII text.

After you enter the complete command syntax, the password is set. You are not prompted to enter the password a second time.

This command enhancement allows you to configure manager, operator, and 802.1X port-access passwords using the CLI in only one step (instead of entering the `password` command and then being prompted twice to enter the actual password).

Release M.10.32 Enhancements

Release M.10.32 includes the following enhancement:

- **Enhancement (PR_1000376626)** — Enhanced the CLI "qos dscp-map he" help and "show dscp-map" text to warn user that inbound classification based on DSCP codepoints only occurs if "qos type-of-service diff-services" is also configured.
- **Enhancement (PR_1000401306)** — Reload "IN/AT" special enhancement.

Scheduled Reload

Additional parameters have been added to the reload command to allow for a scheduled reboot of the switch via the CLI. The scheduled reload feature supports the following capabilities:

It removes the requirement to physically reboot the switch at inconvenient times (for example, at 1:00 in the morning). Instead, a `reload at 1:00 mm/dd` command can be executed (where *mm/dd* is the date the switch is scheduled to reboot). It provides a safety net in situations where a change is made from a remote location to the running config that inadvertently causes loss of management access. For example, a newly configured ACL might deny access to the switch from the management station's IP address such that the telnet session ceases to function. Scheduling a `reload` after command (timed to execute after the necessary configuration work is completed) will ensure that the switch will reboot automatically. Assuming the ACL changes were not saved to the startup config, telnet access will then be restored. If the ACL work is completed successfully, with no loss of access, the scheduled reboot can be cancelled with the `reload cancel` command.

Examples:

To schedule a reload in 15 minutes:

```
ProCurve# reload after 15
```

Enhancements

Release M.10.33 Enhancements

To schedule a reload in 3 hours:

```
ProCurve# reload after 03:00
```

To schedule a reload for the same time the following day:

```
ProCurve# reload after 01:00:00
```

To schedule a reload for the same day at 12:05:

```
ProCurve# reload at 12:05
```

To schedule a reload on some future date:

```
ProCurve# reload at 12:05 01/01/2007
```

Release M.10.33 Enhancements

Release M.10.33 includes the following enhancement:

- **Enhancement (PR_1000408960)** — RADIUS-Assigned GVRP VLANs.

How RADIUS-Based Authentication Affects VLAN Operation

Using a RADIUS server to authenticate clients, you can provide port-level security protection from unauthorized network access for the following authentication methods:

- **802.1X:** Port-based or client-based access control to open a port for client access after authenticating valid user credentials.
- **MAC address:** Authenticates a device's MAC address to grant access to the network.
- **Web-browser interface:** Authenticates clients for network access using a web page for user login.

Note

You can use 802.1X (port-based or client-based) authentication and either Web or MAC authentication at the same time on a port, with a maximum of 32 clients allowed on the port. (The default is one client.) Web authentication and MAC authentication are mutually exclusive on the same port. Also, you must disable LACP on ports configured for any of these authentication methods. For more information, refer to the “Configuring Port-Based and User-Based Access Control (802.1X)” and “Web and MAC Authentication” chapters of the *Access Security Guide*.

VLAN Assignment on a ProCurve Port

Following client authentication, VLAN configurations on a ProCurve port are managed as follows when you use 802.1X, MAC, or Web authentication:

- The port resumes membership in any tagged VLANs for which it is already assigned in the switch configuration. Tagged VLAN membership allows a port to be a member of multiple VLANs simultaneously.

- The port is temporarily assigned as a member of an untagged (static or dynamic) VLAN for use during the client session according to the following order of options.
 - a. The port joins the VLAN to which it has been assigned by a RADIUS server during client authentication.
 - b. If RADIUS authentication does not include assigning the port to a VLAN, then the switch assigns the port to the authorized-client VLAN configured for the authentication method.
 - c. If the port does not have an authorized-client VLAN configured, but is configured for membership in an untagged VLAN, the switch assigns the port to this untagged VLAN.

Operating Notes

- During client authentication, a port assigned to a VLAN by a RADIUS server or an authorized-client VLAN configuration is an untagged member of the VLAN for the duration of the authenticated session. This applies even if the port is also configured in the switch as a tagged member of the same VLAN. The following restrictions apply:
 - If the port is assigned as a member of an untagged *static* VLAN, the VLAN must already be configured on the switch. If the static VLAN configuration does not exist, the authentication fails.
 - If the port is assigned as a member of an untagged *dynamic* VLAN that was learned through GVRP, the dynamic VLAN configuration must exist on the switch at the time of authentication and GVRP-learned dynamic VLANs for port-access authentication must be enabled

If the dynamic VLAN does not exist or if you have not enabled the use of a dynamic VLAN for authentication sessions on the switch, the authentication fails.

- To enable the use of a GVRP-learned (dynamic) VLAN as the untagged VLAN used in an authentication session, enter the **aaa port-access gvrp-vlans** command.
- Enabling the use of dynamic VLANs in an authentication session offers the following benefits:
 - You avoid the need of having static VLANs pre-configured on the switch.
 - You can centralize the administration of user accounts (including user VLAN IDs) on a RADIUS server.

For information on how to enable the switch to dynamically create 802.1Q-compliant VLANs on links to other devices using the GARP VLAN Registration Protocol (GVRP), refer to the “GVRP” chapter in the *Advanced Traffic Management Guide*.

- For an authentication session to proceed, a ProCurve port must be an untagged member of the (static or dynamic) VLAN assigned by the RADIUS server (or an authorized-client VLAN configuration). The port temporarily drops any current untagged VLAN membership.

If the port is not already a member of the RADIUS-assigned (static or dynamic) untagged VLAN, the switch temporarily reassigns the port as an untagged member of the required VLAN (for the duration of the session). *At the same time, if the ProCurve port is already configured as an untagged member of a different VLAN, the port loses access to the other VLAN for the duration of the session.* (A port can be an untagged member of only one VLAN at a time.)

When the authentication session ends, the switch removes the temporary untagged VLAN assignment and re-activates the temporarily disabled, untagged VLAN assignment.

- If GVRP is already enabled on the switch, the temporary untagged (static or dynamic) VLAN created on the port for the authentication session is advertised as an existing VLAN.

If this temporary VLAN assignment causes the switch to disable a different untagged static or dynamic VLAN configured on the port, the disabled VLAN assignment is not advertised. When the authentication session ends, the switch:

- Removes the temporary untagged VLAN assignment and stops advertising it.
- Re-activates and resumes advertising the temporarily disabled, untagged VLAN assignment.
- If you modify a VLAN ID configuration on a port during an 802.1X, MAC, or Web authentication session, the changes do not take effect until the session ends.
- When a switch port is configured with RADIUS-based authentication to accept multiple 802.1X and/or MAC or Web authentication client sessions, all authenticated clients must use the same port-based, untagged VLAN membership assigned for the earliest, currently active client session.

Therefore, on a port where one or more authenticated client sessions are already running, all such clients are on the same untagged VLAN. If a RADIUS server subsequently authenticates a new client, but attempts to re-assign the port to a different, untagged VLAN than the one already in use for the previously existing, authenticated client sessions, the connection for the new client will fail.

Example of Untagged VLAN Assignment in a RADIUS-Based Authentication Session

The following example shows how an untagged static VLAN is temporarily assigned to a port for use during an 802.1X authentication session. In the example, an 802.1X-aware client on port A2 has been authenticated by a RADIUS server for access to VLAN 22. However, port A2 is not configured as a member of VLAN 22 but as a member of untagged VLAN 33 as shown in Figure 8.

```

----- CONSOLE - MANAGER MODE -----
Switch Configuration - VLAN - VLAN Port Assignment

Port  default_vlan  vlan_22  vlan_33  vlan_44
----+-----
A1 | Untagged   Tagged   No       No
A2 | No         No      Untagged No
A3 | Untagged   Forbid  Forbid   Forbid
A4 | Untagged   Tagged  Tagged   Tagged
:   :           :       :       :
:   :           :       :       :
Actions->  Cancel  Edit   Save   Help
Cancel changes and return to previous screen.
Use arrow keys to change action selection and <Enter> to execute
  
```

Scenario: An authorized 802.1X client requires access to VLAN 22 from port A2. However, access to VLAN 22 is blocked (not untagged or tagged) on port A2 and VLAN 33 is untagged on port A2.

Figure 8. Example of an Active VLAN Configuration

In Figure [Figure 8](#), if RADIUS authorizes an 802.1X client on port A2 with the requirement that the client use VLAN 22, then:

- VLAN 22 becomes available as Untagged on port A2 for the duration of the session.
- VLAN 33 becomes unavailable to port A2 for the duration of the session (because there can be only one untagged VLAN on any port).

To view the temporary VLAN assignment as a change in the active configuration, use the **show vlan <vlan-id>** command as shown in Figure [Figure 9](#), where **<vlan-id>** is the (static or dynamic) VLAN used in the authenticated client session.

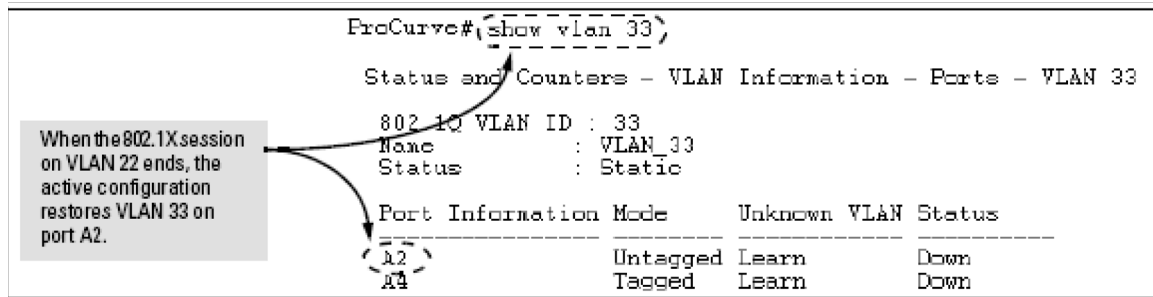


Figure 9. Active Configuration for VLAN 22 Temporarily Changes for the 802.1X Session

However, as shown in Figure [Figure 8](#), because VLAN 33 is configured as untagged on port A2 and because a port can be untagged on only one VLAN, port A2 loses access to VLAN 33 for the duration of the 802.1X session on VLAN 22.

You can verify the temporary loss of access to VLAN 33 by entering the **show vlan 33** command as shown in Figure [Figure 10](#).

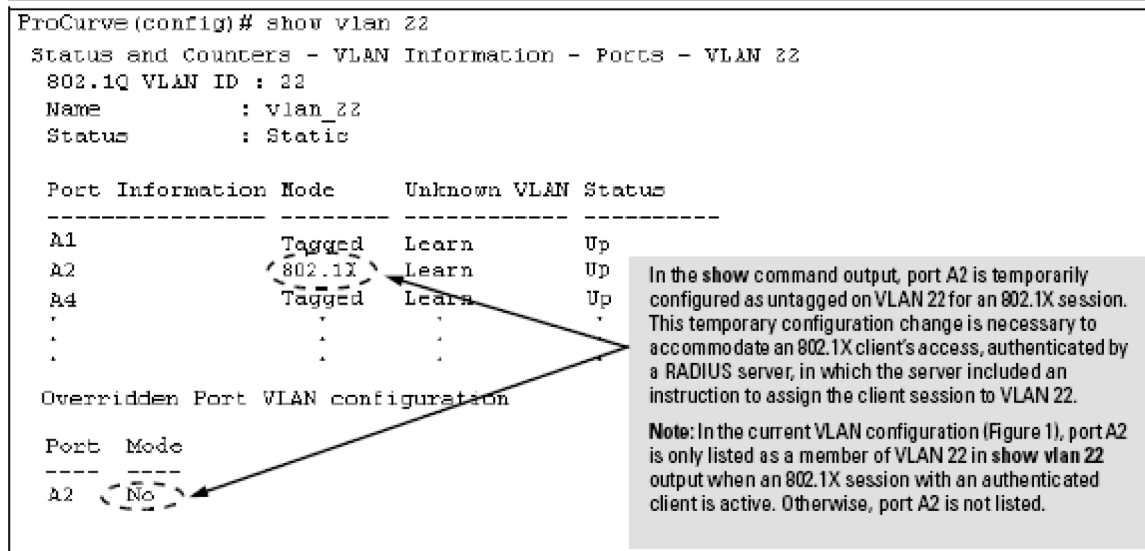


Figure 10. Active Configuration for VLAN 33 Temporarily Drops Port 22 for the 802.1X Session

When the 802.1X client session on port A2 ends, the port removes the temporary untagged VLAN membership. The static VLAN (VLAN 33) that is “permanently” configured as untagged on the port becomes available again. Therefore, when the RADIUS-authenticated 802.1X session on port A2 ends, VLAN 22 access on port A2 also ends, and the untagged VLAN 33 access on port A2 is restored as shown in Figure Figure 11.

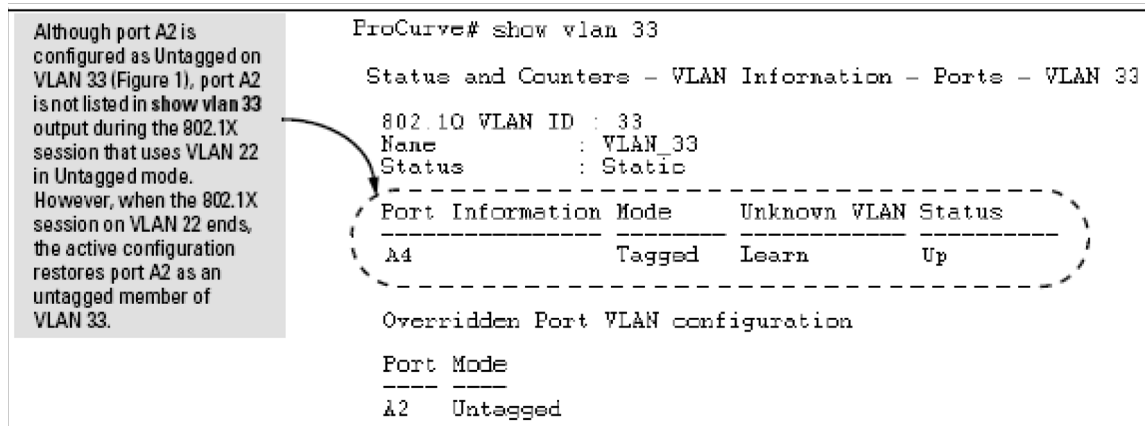


Figure 11. The Active Configuration for VLAN 33 Restores Port A2 After the 802.1X Session Ends

Enabling the Use of GVRP-Learned Dynamic VLANs in Authentication Sessions

Syntax: `aaa port-access gvrp-vlans`

Enables the use of dynamic VLANs (learned through GVRP) in the temporary untagged VLAN assigned by a RADIUS server on an authenticated port in an 802.1X, MAC, or Web authentication session.

*Enter the **no** form of this command to disable the use of GVRP-learned VLANs in an authentication session.*

For information on how to enable a switch to dynamically create 802.1Q-compliant VLANs, refer to the “GVRP” chapter in the Access Security Guide.

Notes:

1. If a port is assigned as a member of an untagged dynamic VLAN, the dynamic VLAN configuration must exist at the time of authentication and GVRP for port-access authentication must be enabled on the switch.

If the dynamic VLAN does not exist or if you have not enabled the use of a dynamic VLAN for authentication sessions on the switch, the authentication fails.

Syntax: `aaa port-access gvrp-vlans`

—Continued—

*2. After you enable dynamic VLAN assignment in an authentication session, it is recommended that you use the **interface unknown-vlans** command on a per-port basis to prevent denial-of-service attacks. The **interface unknown-vlans** command allows you to:*

- *Disable the port from sending advertisements of existing GVRP-created VLANs on the switch.*
- *Drop all GVRP advertisements received on the port.*

For more information, refer to the “GVRP” chapter in the Advanced Traffic Management Guide.

3. If you disable the use of dynamic VLANs in an authentication session using the **no aaa port-access gvrp-vlans** command, client sessions that were authenticated with a dynamic VLAN continue and are not deauthenticated.

(This behavior differs from how static VLAN assignment is handled in an authentication session. If you remove the configuration of the static VLAN used to create a temporary client session, the 802.1X, MAC, or Web authenticated client is deauthenticated.)

However, if a RADIUS-configured dynamic VLAN used for an authentication session is deleted from the switch through normal GVRP operation (for example, if no GVRP advertisements for the VLAN are received on any switch port), authenticated clients using this VLAN are deauthenticated.

For information on how static and dynamic VLANs are assigned in a RADIUS-based 802.1X, MAC, or Web authentication session, refer to the “How RADIUS-Based Authentication Affects VLAN Operation” section in the “RADIUS Authentication and Accounting” chapter of the Access Security Guide.

Release M.10.34 Enhancements

Release M.10.34 includes the following enhancement:

- **Enhancement (PR_1000412747)** — TACACS+ Single Sign-on for Administrators

Concurrent TACAS+ and SFTP

It is now possible to have SFTP/SCP sessions run concurrently with TACACS+ authentication. Because the initial login must be with a username/password that has manager level privileges, you must configure TACACS+ single sign-on in order for TACACS+ and SFTP/SCP to coexist.

To configure TACACS+ single sign-on, use the **aaa authentication login privilege-mode** command.

Syntax: aaa authentication

<login [privilege-mode] >

*Selects the Operator access level. If the **privilege-mode** option is entered, TACACS+ is enabled for a single login. The authorized privilege level (Operator or Manager) is granted by the TACACS+ server.*

Default: Single login disabled.

Release M.10.35 Enhancements

Release M.10.35 includes the following enhancement:

- **Enhancement (PR_1000419928)** — The Dynamic ARP Protection feature was added.

Dynamic ARP Protection

Introduction

On the VLAN interfaces of a routing switch, dynamic ARP protection ensures that only valid ARP requests and responses are relayed or used to update the local ARP cache. ARP packets with invalid IP-to-MAC address bindings advertised in the source protocol address and source physical address fields are discarded. For more information about the ARP cache, refer to “ARP Cache Table” in the *Multicast and Routing Guide*.

ARP requests are ordinarily broadcast and received by all devices in a broadcast domain. Most ARP devices update their IP-to-MAC address entries each time they receive an ARP packet even if they did not request the information. This behavior makes an ARP cache vulnerable to attacks.

Because ARP allows a node to update its cache entries on other systems by broadcasting or unicasting a gratuitous ARP reply, an attacker can send his own IP-to-MAC address binding in the reply that causes all traffic destined for a VLAN node to be sent to the attacker's MAC address. As a result, the attacker can intercept traffic for other hosts in a classic "man-in-the-middle" attack. The attacker gains access to any traffic sent to the poisoned address and can capture passwords, e-mail, and VoIP calls or even modify traffic before resending it.

Another way in which the ARP cache of known IP addresses and associated MAC addresses can be poisoned is through unsolicited ARP responses. For example, an attacker can associate the IP address of the network gateway with the MAC address of a network node. In this way, all outgoing traffic is prevented from leaving the network because the node does not have access to outside networks. As a result, the node is overwhelmed by outgoing traffic destined to another network.

Dynamic ARP protection is designed to protect your network against ARP poisoning attacks in the following ways:

- Allows you to differentiate between trusted and untrusted ports.
- Intercepts all ARP requests and responses on untrusted ports before forwarding them.
- Verifies IP-to-MAC address bindings on untrusted ports with the information stored in the lease database maintained by DHCP snooping and user-configured static bindings (in non-DHCP environments):
 - If a binding is valid, the switch updates its local ARP cache and forwards the packet.

- If a binding is invalid, the switch drops the packet, preventing other network devices from receiving the invalid IP-to-MAC information.

DHCP snooping intercepts and examines DHCP packets received on switch ports before forwarding the packets. DHCP packets are checked against a database of DHCP binding information. Each binding consists of a client MAC address, port number, VLAN identifier, leased IP address, and lease time. The DHCP binding database is used to validate packets by other security features on the switch.

If you have already enabled DHCP snooping on a switch, you may also want to add static IP-to-MAC address bindings to the DHCP snooping database so that ARP packets from devices that have been assigned static IP addresses are also verified.

- Supports additional checks to verify source MAC address, destination MAC address, and IP address.

ARP packets that contain invalid IP addresses or MAC addresses in their body that do not match the addresses in the Ethernet header are dropped.

When dynamic ARP protection is enabled, only ARP request and reply packets with valid IP-to-MAC address bindings in their packet header are relayed and used to update the ARP cache.

Dynamic ARP protection is implemented in the following ways on a switch:

- You can configure dynamic ARP protection only from the CLI; you cannot configure this feature from the web or menu interfaces.
- Line rate—Dynamic ARP protection copies ARP packets to the switch CPU, evaluates the packets, and then re-forwards them through the switch software. During this process, if ARP packets are received at too high a line rate, some ARP packets may be dropped and will need to be retransmitted.
- The SNMP MIB, HP-ICF-ARP-PROTECT-MIB, is created to configure dynamic ARP protection and to report ARP packet-forwarding status and counters.

Enabling Dynamic ARP Protection

To enable dynamic ARP protection for VLAN traffic on a routing switch, enter the **arp protect vlan** command at the global configuration level.

Syntax: [no] arp protect vlan [*vlan-range*]

vlan-range *Specifies a VLAN ID or a range of VLAN IDs from one to 4094; for example, 1–200.*

An example of the **arp protect vlan** command is shown here:

```
ProCurve(config)# arp protect vlan 1-101
```


Configuring Trusted Ports

In a similar way to DHCP snooping, dynamic ARP protection allows you to configure VLAN interfaces in two categories: trusted and untrusted ports. ARP packets received on trusted ports are forwarded without validation.

By default, all ports on a switch are untrusted. If a VLAN interface is untrusted:

- The switch intercepts all ARP requests and responses on the port.
- Each intercepted packet is checked to see if its IP-to-MAC binding is valid. If a binding is invalid, the switch drops the packet.

You must configure trusted ports carefully. For example, in the topology in Figure 8, Switch B may not see the leased IP address that Host 1 receives from the DHCP server. If the port on Switch B that is connected to Switch A is untrusted and if Switch B has dynamic ARP protection enabled, it will see ARP packets from Host 1 as invalid, resulting in a loss of connectivity.

On the other hand, if Switch A does not support dynamic ARP protection and you configure the port on Switch B connected to Switch A as trusted, Switch B opens itself to possible ARP poisoning from hosts attached to Switch A.




Figure 12. Configuring Trusted Ports for Dynamic ARP Protection

Take into account the following configuration guidelines when you use dynamic ARP protection in your network:

- You should configure ports connected to other switches in the network as trusted ports. In this way, all network switches can exchange ARP packets and update their ARP caches with valid information.
- Switches that do not support dynamic ARP protection should be separated by a router in their own Layer 2 domain. Because ARP packets do not cross Layer 2 domains, the unprotected switches cannot unknowingly accept ARP packets from an attacker and forward them to protected switches through trusted ports.

To configure one or more Ethernet interfaces that handle VLAN traffic as trusted ports, enter the **arp protect trust** command at the global configuration level. The switch does not check ARP requests and responses received on a trusted port.

Syntax: [no] arp protect trust <port-list>

port-list *Specifies a port number or a range of port numbers. Separate individual port numbers or ranges of port numbers with a comma; for example: c1-c3, c6.*

An example of the **arp protect trust** command is shown here:

```
ProCurve(config)# arp protect trust b1-b4, d1
```

Adding an IP-to-MAC Binding to the DHCP Database

A routing switch maintains a DHCP binding database, which is used for DHCP and ARP packet validation. Both the DHCP snooping and DHCP Option 82 insertion features maintain the lease database by learning the IP-to-MAC bindings on untrusted ports. Each binding consists of the client MAC address, port number, VLAN identifier, leased IP address, and lease time.

If your network does not use DHCP or if some network devices have fixed, user-configured IP addresses, you can enter static IP-to-MAC address bindings in the DHCP binding database. The switch uses manually configured static bindings for DHCP snooping and dynamic ARP protection.

To add the static configuration of an IP-to-MAC binding for a port to the database, enter the **ip source binding** command at the global configuration level.

Syntax: [no] ip source binding <mac-address> vlan <vlan-id> <ip-address>
interface <port-number>

mac-address *Specifies a MAC address to bind with a VLAN and IP address on the specified port in the DHCP binding database.*

vlan <vlan-id> *Specifies a VLAN ID number to bind with the specified MAC and IP addresses on the specified port in the DHCP binding database.*

ip-address *Specifies an IP address to bind with a VLAN and MAC address on the specified port in the DHCP binding database.*

interface <port-number> *Specifies the port number on which the IP-to-MAC address and VLAN binding is configured in the DHCP binding database.*

An example of the **ip source binding** command is shown here:

```
ProCurve(config)# ip source binding 0030c1-7f49c0  
interface vlan 100 10.10.20.1 interface A4
```

Note

Note that the **ip source binding** command is the same command used by the Dynamic IP Lockdown feature to configure static bindings. The Dynamic ARP Protection and Dynamic IP Lockdown features share a common list of source IP-to-MAC bindings.

Configuring Additional Validation Checks on ARP Packets

Dynamic ARP protection can be configured to perform additional validation checks on ARP packets. By default, no additional checks are performed. To configure additional validation checks, enter the **arp protect validate** command at the global configuration level.

Syntax: [no] arp protect validate <[src-mac] | [dst-mac] | [ip]>

- src-mac** *(Optional) Drops any ARP request or response packet in which the source MAC address in the Ethernet header does not match the sender MAC address in the body of the ARP packet.*
- dst-mac** *(Optional) Drops any unicast ARP response packet in which the destination MAC address in the Ethernet header does not match the target MAC address in the body of the ARP packet.*
- ip** *(Optional) Drops any ARP packet in which the sender IP address is invalid. Drops any ARP response packet in which the target IP address is invalid. Invalid IP addresses include: 0.0.0.0, 255.255.255.255, all IP multicast addresses, and all Class E IP addresses.*

You can configure one or more of the validation checks. The following example of the **arp protect validate** command shows how to configure the validation checks for source MAC address and destination MAC address:

```
ProCurve(config)# arp protect validate src-mac dst-mac
```

Verifying the Configuration of Dynamic ARP Protection

To display the current configuration of dynamic ARP protection, including the additional validation checks and the trusted ports that are configured, enter the **show arp protect** command:

```
ProCurve(config)# show arp protect

ARP Protection Information

Enabled Vlans   : 1-4094
Validate       : dst-mac, src-mac

Port   Trust
-----
B1     Yes
B2     Yes
B3     No
B4     No
B5     No
```

Figure 13. The show arp protect Command

Displaying ARP Packet Statistics

To display statistics about forwarded ARP packets, dropped ARP packets, MAC validation failure, and IP validation failures, enter the **show arp protect statistics** command:

```
ProCurve(config)# show arp protect statistics

Status and Counters - ARP Protection Counters for VLAN 1

Forwarded pkts   : 10      Bad source mac      : 2
Bad bindings     : 1       Bad destination mac: 1
Malformed pkts  : 0       Bad IP address      : 0

Status and Counters - ARP Protection Counters for VLAN 2

Forwarded pkts   : 1       Bad source mac      : 1
Bad bindings     : 1       Bad destination mac: 1
Malformed pkts  : 1       Bad IP address      : 1
```

Figure 14. Show arp protect statistics Command

Monitoring Dynamic ARP Protection

When dynamic ARP protection is enabled, you can monitor and troubleshoot the validation of ARP packets with the **debug arp protect** command. Use this command when you want to debug the following conditions:

- The switch is dropping valid ARP packets that should be allowed.
- The switch is allowing invalid ARP packets that should be dropped.

```
ProCurve(config)# debug arp protect

1. ARP request is valid
"DARPP: Allow ARP request 000000-000001,10.0.0.1 for 10.0.0.2 port A1,
vlan "

2. ARP request detected with an invalid binding
"DARPP: Deny ARP request 000000-000003,10.0.0.1 port A1, vlan 1"

3. ARP response with a valid binding
"DARPP: Allow ARP reply 000000-000002,10.0.0.2 port A2, vlan 1"

4. ARP response detected with an invalid binding
"DARPP: Deny ARP reply 000000-000003,10.0.0.2 port A2, vlan 1"
```

Figure 15. Example of debug arp protect Command

Release M.10.36 Enhancements

Software fixes only, no new enhancements.

Release M.10.37 Enhancements

Release M.10.37 includes the following enhancement:

- **Enhancement (PR_1000369492)** — Update of the MSTP implementation to the latest IEEE P802.1Q-REV/D5.0 specifications to stay in sync with the protocol evolution.

For more information on selected configuration options and updated MSTP port parameters, see [“Configuring MSTP Port Connectivity Parameters”](#) below.

Configuring MSTP Port Connectivity Parameters

With release K.12.04, all ports are configured as auto-edge-ports by default, and the spanning tree **edge-port** option has been removed. This section describes selected **spanning-tree <port-list>** command parameters for enhanced operation.

Basic port connectivity parameters affect spanning-tree links at the global level. Therefore, in most cases, ProCurve recommends that you use the revised default settings for these parameters and apply changes on a per-port basis only where a non-default setting is clearly indicated by the circumstances of individual links (for example, see the **root-guard** option below).

To display the spanning-tree settings for each port, use the **show spanning-tree config** command.

Syntax: [no] spanning-tree < port-list > < auto-edge-port | admin-edge-port | mcheck | root-guard | tcn-guard >

[auto-edge-port]

*Enables **auto-edge-port** operation for MSTP, and supports the automatic detection of edge ports. (Default: **Yes**, enabled)*

*The port will look for BPDUs for 3 seconds; if there are none it begins forwarding packets. If **admin-edge-port** is enabled for a port, the setting for **auto-edge-port** is ignored whether set to yes or no. If **admin-edge-port** is disabled, and **auto-edge-port** has not been disabled, then the **auto-edge-port** setting controls the behavior of the port.*

*The **no spanning-tree < port-list > auto-edge-port** command disables **auto-edge-port** operation on the specified ports.*

[admin-edge-port]

*Enables **admin-edge-port** for RSTP/MSTP. If a bridge or switch is detected on the segment, the port automatically operates as non-edge, not enabled. (Default: **No** - disabled)*

*If **admin-edge-port** is disabled on a port and **auto-edge-port** has not been disabled, the **auto-edge-port** setting controls the behavior of the port.*

*The **no spanning-tree < port-list > admin-edge-port** command disables **admin-edge-port** operation on the specified ports.*

[mcheck]

Forces a port to send RSTP/MSTP BPDUs for 3 seconds. This allows for another switch connected to the port and running RSTP to establish its connection quickly and for identifying switches running 802.1D STP. If the whole-switch force-version parameter is set to stp-compatible, the switch ignores the mcheck setting and sends 802.1D STP BPDUs out all ports.

[root-guard]

*MSTP only. When a port is enabled as **root-guard**, it cannot be selected as the root port even if it receives superior STP BPDUs. The port is assigned an “alternate” port role and enters a blocking state if it receives superior STP BPDUs. The BPDUs received on a **root-guard** port are ignored. All other BPDUs are accepted and the external devices may belong to the spanning tree as long as they do not claim to be the Root device. (Default: **No** - disabled)*

***Note:** In standard Spanning Tree Protocol operation, the calculation of active network topologies may be an issue when switches outside the core region of a network are under shared or limited administrative control. Such a switch may become a Root Bridge for the entire network and create non-optimal forwarding paths. By enabling the **root-guard** feature on ports that face outside the core network, external boundaries for the core network are created to ensure the Root Bridge is located within the core network.*

[tcn-guard]

*When **tcn-guard** is enabled for a port, it causes the port to stop propagating received topology change notifications and topology changes to other ports. (Default: **No** - disabled)*

Syntax: spanning-tree < port-list > < hello-time | path-cost | point-to-point-mac | priority >

[hello-time < global | 1 - 10 >

*When the switch is the CIST root, this parameter specifies the interval (in seconds) between periodic BPDU transmissions by the designated ports. This interval also applies to all ports in all switches downstream from each port in the < port-list >. A setting of **global** indicates that the ports in < port-list > on the CIST root are using the value set by the global spanning-tree **hello-time** value. When a given switch “X” is not the CIST root, the per-port **hello-time** for all active ports on switch “X” is propagated from the CIST root, and is the same as the **hello-time** in use on the CIST root port in the currently active path from switch “X” to the CIST root. (That is, when switch “X” is not the CIST root, then the upstream CIST root’s port **hello-time** setting overrides the **hello-time** setting configured on switch “X”. (Default Per-Port setting: **Use Global**. Default Global Hello-Time: **2**.)*

[path-cost < auto | 1..200000000 >]

Assigns an individual port cost that the switch uses to determine which ports are forwarding ports in a given spanning tree. In the default configuration (auto) the switch determines a port’s path cost by the port’s type:

- 10 Mbps: **2000000***
- 100 Mbps: **200000***
- 1 Gbps: **20000***

point-to-point-mac <true | false | auto >

This parameter informs the switch of the type of device to which a specific port connects.

True (default): Indicates a point-to-point link to a device such as a switch, bridge, or end-node.

False: Indicates a connection to a hub (which is a shared LAN segment).

Auto: Causes the switch to set False on the port if it is not running at full duplex. (Connections to hubs are half-duplex.)

priority < 0..15 >

MSTP uses this parameter to determine the port(s) to use for forwarding. The port with the lowest assigned value has the highest priority. While the actual priority range is 0 to 240, this command specifies the priority as a multiplier (0-15) of 16. That is, when you specify a priority multiplier of 0-15, the actual priority assigned to the switch is:

$$(priority-multiplier) \times 16 = priority$$

The default priority-multiplier value is 8.

*For example, if you configure “2” as the priority multiplier for a given port, then the actual priority is 32. Thus, after you specify the port priority multiplier, the switch displays the actual port priority (and not the multiplier) in the **show spanning-tree config** display. You can view the actual multiplier setting for ports by executing **show running** and looking for an entry in this form:*

```
spanning-tree <port-list> priority <priority-multiplier>
```

*For example, configuring port 2 with a priority multiplier of “3” results in this line in the **show running-config** output:*

```
spanning-tree B2 priority 3
```

Release M.10.38 Enhancements

Release M.10.38 includes the following enhancement:

- **Enhancement (PR_1000428642)** — SNMP v2c describes two different notification-type PDUs: traps and informs. Prior to this software release, only the traps sub-type was supported. This enhancement adds support for informs.

Send SNMP v2c Informs

Enabling and Configuring SNMP Informs

You can use the **snmp-server informs** command (SNMPv2c and SNMPv3 versions) to send notifications when certain events occur. When an SNMP Manager receives an informs request, it can send an SNMP response back to the sending agent. This lets the agent know that the informs request reached its destination and that traps can be sent successfully to that destination.

Informs requests can be sent several times until a response is received from the SNMP manager or the configured retry limits are reached. The request may also timeout.

To enable SNMP informs, enter this command:

Syntax: [no] snmp-server enable informs

Enables or disables the informs option for SNMP.

Default: Disabled

To configure SNMP informs request options, use the following commands.

Syntax: [no] snmp-server informs [retries<retries>] [timeout<seconds>] [pending <pending>]

Allows you to configure options for SNMP informs requests.

retries: Maximum number of times to resend an informs request. Default: 3

timeout: Number of seconds to wait for an acknowledgement before resending the informs request. Default: 30 seconds

pending: *Maximum number of informs waiting for acknowledgement at any one time. When the maximum configured number is reached, older pending informs are discarded. Default: 25*

To specify the manager that receives the informs request, use the **snmp-server host** command.

Syntax: snmp-server host < ip-address >[<traps | informs>] [version <1 | 2c | 3>]< community-string >

Using community name and destination IP address, this command designates a destination network-management station for receiving SNMP event log messages from the switch. If you do not specify the event level, then the switch does not send event log messages as traps. You can specify up to 10 trap receivers (network management stations).

Note: *In all cases, the switch sends any threshold trap(s) or informs to the network management station(s) that explicitly set the threshold(s).*

[traps | informs>]

Select whether SNMP traps or informs are sent to this management station. For more information on SNMP informs, see [“Enabling and Configuring SNMP Informs” on page 119](#).

[version <1 | 2c | 3>]

Select the version of SNMP being used.

Note: SNMP informs are supported on version 2c or 3 only.

[<none | all | non-info | critical | debug>]

Options for sending switch Event Log messages to a trap receiver. The levels specified with these options apply only to Event Log messages, and not to threshold traps.

You can see if informs are enabled or disabled with the **show snmp-server** command as shown in Figure 11.

```
ProCurve(config)# show snmp-server
SNMP Communities
  Community Name      MIB View Write Access
  -----
  public              Manager  Unrestricted
Trap Receivers
  Link-Change Traps Enabled on Ports [All] : All
  Send Authentication Traps [No] : No
  [ Informs [Yes] : Yes ]
  Address              | Community      Events Sent in Trap
  -----
--
Excluded MIBs

Snmpp Response Pdu Source-IP Information
  Selection Policy    : Default rfc1517
Trap Pdu Source-IP Information
  Selection Policy    : Default rfc1517
```

Figure 11. Example Showing SNMP Informs Option Enabled

Release M.10.39 Enhancements

Release M.10.39 includes the following enhancement:

- **Enhancement (PR_1000428213)** — This software enhancement adds the ability to configure a secondary authentication method to be used when the RADIUS server is unavailable for the primary port-access method.

RADIUS Server Unavailable

Overview

In certain situations, RADIUS servers can become isolated from the network. Users are not able to access the network resources configured with RADIUS access protection and are rejected. To address this situation, configuring the “**authorized**” secondary authentication method allows users unconditional access to the network when the primary authentication method fails because the RADIUS servers are unreachable.

Configuring RADIUS Authentication

You can configure the switch for RADIUS authentication through the following access methods:

- **Console:** Either direct serial-port connection or modem connection.
- **Telnet:** Inbound Telnet must be enabled (the default).
- **SSH:** To use RADIUS for SSH access, first configure the switch for SSH operation.
- **Web:** Enables RADIUS authentication for web browser interface access to the switch.

You can configure **radius** as the primary password authentication method for the above access methods. You also need to select either **local**, **none**, or **authorized** as a secondary, or backup, method..

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

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

`[< local | none | authorized >]`

Provides options for secondary authentication (default: none).

Caution

Configuring **authorized** as the secondary authentication method used when there is a failure accessing the RADIUS servers allows clients to access the network unconditionally. Use this method with care.

You can configure **local**, **chap-radius** or **eap-radius** as the primary password authentication method for the port-access method. You also need to select **none** or **authorized** as a secondary, or backup, method.

Syntax: aaa authentication port-access <chap-radius leap-radius | local>

*Configures **local**, **chap-radius**, or **eap-radius** as the primary password authentication method for port-access. The default primary authentication is **local**.*

[<none | authorized >]

*Provides options for secondary authentication. The **none** option specifies that a backup authentication method is not used. The **authorized** option allows access without authentication. (default: **none**).*

You can configure **chap-radius** as the primary password authentication method for web-based or mac-based port-access methods. You also need to select **none** or **authorized** as a secondary, or backup, method.

Syntax: aaa authentication <mac-based | web-based> chap-radius

*Configures **chap-radius** as the primary password authentication method for mac-based or web-based port access.*

[<none | authorized >]

*Provides options for secondary authentication. The **none** option specifies that a backup authentication method is not used. The **authorized** option allows access without authentication. (default: **none**).*

Figure 1 shows an example of the **show authentication** command displaying **authorized** as the secondary authentication method for port-access, Web-auth access, and Mac-auth access. Since the configuration of **authorized** means no authentication will be performed and the client has unconditional access to the network, the “Enable Primary” and “Enable Secondary” fields are not applicable (N/A).

```
ProCurve(config)# show authentication

Status and Counters - Authentication Information

Login Attempts : 3
Respect Privilege : Disabled
```

	Login	Login	Enable	Enable	
Access Task	Primary	Secondary	Primary	Secondary	
Console	Local	None	Local	None	<div style="background-color: #e0e0e0; padding: 5px;"> The access methods with secondary authentication configured as authorized allows the client access to the network even if the RADIUS server is unreachable. </div>
Telnet	Local	None	Local	None	
Port-Access	Local	Authorized	N/A	N/A	
Webui	Local	None	Local	None	
SSH	Local	None	Local	None	
Web-Auth	ChapRadius	Authorized	N/A	N/A	
MAC-Auth	ChapRadius	Authorized	N/A	N/A	

Figure 12. Example of AAA Authentication Using Authorized for the Secondary Authentication Method

Specifying the MAC Address Format

The MAC address format command has been enhanced to allow upper-case letters to be used for the hexadecimal numbers when indicating the MAC address in RADIUS packets for MAC-based authentication.

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

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

- no-delimiter** — specifies an aabbccddeeff format.
- single-dash** — specifies an aabbcc-ddeeff format.
- multi-dash** — specifies an aa-bb-cc-dd-ee-ff format.
- multi-colon** — specifies an aa:bb:cc:dd:ee:ff format.
- no-delimiter-uppercase** — specifies an AABBCCDDEEFF format.
- single-dash-uppercase** — specifies an AABBCD-DDEEFF format
- multi-dash-uppercase** — specifies an AA-BB-CC-DD-EE-FF format
- multi-colon-uppercase** — specifies an AA:BB:CC:DD:EE:FF format.

For example, using the multi-colon-uppercase option, the MAC address would appear as follows:

AA:BB:CC:DD:EE:FF

- **Enhancement (PR_1000415155)** — The ARP age timer was enhanced from the previous limit of 240 minutes to allow for configuration of values up to 1440 minutes (24 hours) or "infinite" (99,999,999 seconds or 3.2 years).

ARP Age Timer Increase

The ARP age is the amount of time the switch keeps a MAC address learned through ARP in the ARP cache. The switch resets the timer to zero each time the ARP entry is refreshed and removes the entry if the timer reaches the ARP age.

You can increase the ARP age timeout maximum to 24 hours or more with this command:

Syntax: [no] ip arp-age <[1...1440] | infinite>

Allows the ARP age to be set from 1 to 1440 minutes (24 hours). If the option "infinite" is configured, the internal ARP age timeout is set to 99,999,999 seconds (approximately 3.2 years). An arp-age value of 0 (zero) is stored in the configuration file to indicate that "infinite" has been configured. This value also displays with the show commands and in the menu display (Menu > Switch Configuration > IP Config).

Default: 20 minutes.

```
ProCurve(config)# ip arp-age 1000
```

Figure 13. Example of Setting the ARP Age Timeout to 1000 Minutes

To view the value of Arp Age timer, enter the **show ip** command as shown in [Figure 14](#).

```
ProCurve(config)# show ip

Internet (IP) Service

  IP Routing : Disabled

  Default Gateway : 15.255.120.1
  Default TTL    : 64
  Arp Age       : 1000
  Domain Suffix :
  DNS server    :

VLAN          | IP Config | IP Address | Subnet Mask | Proxy ARP
-----+-----
DEFAULT_VLAN | Manual    | 15.255.111.13 | 255.255.248.0 | No
```

Figure 14. Example of show ip Command Displaying Arp Age

You can also view the value of the Arp Age timer in the configuration file.

```
ProCurve(config)# show running-config

Running configuration:

; J9091A Configuration Editor; Created on release #K.12.XX

hostname "8200LP"
module 2 type J8702A
module 3 type J8702A
module 4 type J8702A
ip default-gateway 15.255.120.1
[ip_arp_age_1000_]
snmp-server community "public" Unrestricted
snmp-server host 16.180.1.240 "public"
vlan 1
    name "DEFAULT_VLAN"
    untagged B1-B24,C1-C24,D1-D24
    ip address 15.255.120.85 255.255.248.0
    exit
gvrp
spanning-tree
```

Figure 15. Example Showing ip arp-age Value in the Running Config File

You can set or display the **arp-age** value using the menu interface (**Menu > Switch Configuration > IP Config**).

```
ProCurve                                     12-June-2007 14:45:31
=====-- TELNET - MANAGER MODE =====
                Switch Configuration - Internet (IP) Service

IP Routing : Disabled

Default Gateway : 15.255.120.1
Default TTL    : 64
Arp Age       : 1000

IP Config [Manual] : Manual

IP Address    : 15.255.111.11
Subnet Mask   : 255.255.248.0

Actions->   Cancel      Edit      Save      Help
```

Figure 16. Example of the Menu Interface Displaying the Arp Age Value

Enhancements

Release M.10.40 Enhancements

If the ARP cache should become full because entries are not cleared (due to increased timeout limits) you can use the **clear arp** command to remove all non-permanent entries in the ARP cache.

To remove a specific entry in the ARP cache, enter this command:

Syntax: [no] arp IP-ADDRESS

Allows removal of any dynamic entry in the ARP cache.

Release M.10.40 Enhancements

Software fixes only, no new enhancements.

Release M.10.41 Enhancements

Software fixes only, no new enhancements.

Release M.10.42 Enhancements

Software fixes only, no new enhancements.

Release M.10.43 Enhancements

Release M.10.43 includes the following enhancements:

- **Enhancement (PR_1000428642)** — SNMP v2c describes two different notification-type PDUs: traps and informs. Prior to this software release, only the traps sub-type was supported. This enhancement adds support for informs.
- **Enhancement (PR_1000452407)** — The Dynamic IP Lockdown feature was added for the 3400cl series switches.

Dynamic IP Lockdown

The Dynamic IP Lockdown feature is used to prevent IP source address spoofing on a per-port and per-VLAN basis. When dynamic IP lockdown is enabled, IP packets in VLAN traffic received on a port are forwarded only if they contain a known source IP address and MAC address binding for the port. The IP-to-MAC address binding can either be statically configured or learned by the DHCP Snooping feature.

Protection Against IP Source Address Spoofing

Many network attacks occur when an attacker injects packets with forged IP source addresses into the network. Also, some network services use the IP source address as a component in their authentication schemes. For example, the BSD “r” protocols (rlogin, rcp, rsh) rely on the IP source address for packet authentication. SNMPv1 and SNMPv2c also frequently use authorized IP address lists to limit management access. An attacker that is able to send traffic that appears to originate from an authorized IP source address may gain access to network services for which he is not authorized.

Dynamic IP lockdown provides protection against IP source address spoofing by means of IP-level port security. IP packets received on a port enabled for dynamic IP lockdown are only forwarded if they contain a known IP source address and MAC address binding for the port.

Dynamic IP lockdown uses information collected in the DHCP Snooping lease database and through statically configured IP source bindings to create internal, per-port lists. The internal lists are dynamically created from known IP-to-MAC address bindings to filter VLAN traffic on both the source IP address and source MAC address.

Differences Between Switch Platforms

There are some differences in the feature set and operation of Dynamic IP Lockdown, depending on the switch on which it is implemented. These are listed below.

- There is no restriction on GVRP on 3500/5400 switches. On 2600/2800/3400cl switches, Dynamic IP Lockdown is not supported if GVRP is enabled on the switch.
- Dynamic IP Lockdown has the host limits shown in the table below. There is a DHCP snooping limit of 8,000 entries.

Switch	Number of Hosts	Comments
3500/5400	64 bindings per port Up to 4096 bindings per switch	This limit is shared with DHCP snooping because they both use the snooping database.
3400cl/2800	32 bindings per port Up to 32 VLANs with DHCP snooping enabled	This is not guaranteed as the hardware resources are shared with QoS.
2600	8 bindings per port Up to 8 VLANs with DHCP snooping enabled	This is not guaranteed as the hardware resources are shared with QoS.

- A source is considered “trusted” for all VLANs if it is seen on any VLAN without DHCP snooping enabled.
- On the ProCurve switch series 5400 and 3500, dynamic IP lockdown is supported on a port configured for statically configured port-based ACLs.

Prerequisite: DHCP Snooping

Dynamic IP lockdown requires that you enable DHCP snooping as a prerequisite for its operation on ports and VLAN traffic:

- Dynamic IP lockdown only enables traffic for clients whose leased IP addresses are already stored in the lease database created by DHCP snooping or added through a static configuration of an IP-to-MAC binding.

Therefore, if you enable DHCP snooping after dynamic IP lockdown is enabled, clients with an existing DHCP-assigned address must either request a new leased IP address or renew their existing DHCP-assigned address. Otherwise, a client's leased IP address is not contained in the DHCP binding database. As a result, dynamic IP lockdown will not allow inbound traffic from the client.

- It is recommended that you enable DHCP snooping a week before you enable dynamic IP lockdown to allow the DHCP binding database to learn clients' leased IP addresses. You must also ensure that the lease time for the information in the DHCP binding database lasts more than a week.

Alternatively, you can configure a DHCP server to re-allocate IP addresses to DHCP clients. In this way, you repopulate the lease database with current IP-to-MAC bindings.

- The DHCP binding database allows VLANs enabled for DHCP snooping to be known on ports configured for dynamic IP lockdown. As new IP-to-MAC address and VLAN bindings are learned, a corresponding permit rule is dynamically created and applied to the port (preceding the final deny any vlan <VLAN_IDS> rule as shown in the example in Figure 3). These VLAN_IDS correspond to the subset of configured and enabled VLANs for which DHCP snooping has been configured.
- For dynamic IP lockdown to work, a port must be a member of at least one VLAN that has DHCP snooping enabled.
- Disabling DHCP snooping on a VLAN causes Dynamic IP bindings on Dynamic IP Lockdown-enabled ports in this VLAN to be removed. The port reverts back to switching traffic as usual.

Filtering IP and MAC Addresses Per-Port and Per-VLAN

This section contains an example that shows the following aspects of the Dynamic IP Lockdown feature:

- Internal Dynamic IP lockdown bindings dynamically applied on a per-port basis from information in the DHCP Snooping lease database and statically configured IP-to-MAC address bindings
- Packet filtering using source IP address, source MAC address, and source VLAN as criteria

In this example, the following DHCP leases have been learned by DHCP snooping on port 5. VLANs 2 and 5 are enabled for DHCP snooping.

IP Address	MAC Address	VLAN ID
10.0.8.5	001122-334455	2
10.0.8.7	001122-334477	2
10.0.10.3	001122-334433	5

Figure 17. Sample DHCP Snooping Entries

The following example shows an IP-to-MAC address and VLAN binding that have been statically configured in the lease database on port 5.

IP Address	MAC Address	VLAN ID
10.0.10.1	001122-110011	5

Figure 18. An Example of a Static Configuration Entry

Assuming that DHCP snooping is enabled and that port 5 is untrusted, dynamic IP lockdown applies the following dynamic VLAN filtering on port 5:

```
permit 10.0.8.5 001122-334455 vlan 2
permit 10.0.8.7 001122-334477 vlan 2
permit 10.0.10.3 001122-334433 vlan 5
permit 10.0.10.1 001122-110011 vlan 5
deny any vlan 1-10
permit any
```

Figure 19. Example of Internal Statements used by Dynamic IP Lockdown

Note that the **deny any** statement is applied only to VLANs for which DHCP snooping is enabled. The **permit any** statement is applied only to all other VLANs.

Enabling Dynamic IP Lockdown

To enable dynamic IP lockdown on all ports or specified ports, enter the **ip source-lockdown** command at the global configuration level. Use the no form of the command to disable dynamic IP lockdown.

Syntax: [no] ip source-lockdown [port-list]

Enables dynamic IP lockdown globally on all ports or on specified ports on the routing switch.

Operating Notes

- Dynamic IP lockdown is enabled at the port configuration level and applies to all bridged or routed IP packets entering the switch. The only IP packets that are exempt from dynamic IP lockdown are broadcast DHCP request packets, which are handled by DHCP snooping.
- DHCP snooping is a prerequisite for Dynamic IP Lockdown operation. The following restrictions apply:
 - DHCP snooping is required for dynamic IP lockdown to operate. To enable DHCP snooping, enter the **DHCP-Snooping** command at the global configuration level.
 - Dynamic IP lockdown only filters packets in VLANs that are enabled for DHCP snooping. In order for Dynamic IP lockdown to work on a port, the port must be configured for at least one VLAN that is enabled for DHCP snooping.

To enable DHCP snooping on a VLAN, enter the **dhcp-snooping vlan [vlan-id-range]** command at the global configuration level or the **dhcp-snooping** command at the VLAN configuration level.

- Dynamic IP lockdown is not supported on a trusted port. (However, note that the DHCP server must be connected to a trusted port when DHCP snooping is enabled.)

By default, all ports are untrusted. To remove the trusted configuration from a port, enter the **no dhcp-snooping trust <port-list>** command at the global configuration level.

For more information on how to configure and use DHCP snooping, refer to the “Configuring Advanced Threat Protection” chapter in the *Access Security Guide*.

- After you enter the **ip source-lockdown** command (enabled globally with the desired ports entered in *<port-list>*), the dynamic IP lockdown feature remains disabled on a port if any of the following conditions exist:
 - If DHCP snooping has not been globally enabled on the switch.
 - If the port is not a member of at least one VLAN that is enabled for DHCP snooping.
 - If the port is configured as a trusted port for DHCP snooping.

Dynamic IP lockdown is activated on the port only after you make the following configuration changes:

- Enable DHCP snooping on the switch.
- Configure the port as a member of a VLAN that has DHCP snooping enabled.

- Remove the trusted-port configuration.
- You can configure dynamic IP lockdown only from the CLI; this feature cannot be configured from the Web management or menu interface.
- If you enable dynamic IP lockdown on a port, you cannot add the port to a trunk.
- Dynamic IP lockdown must be removed from a trunk before the trunk is removed.

Adding an IP-to-MAC Binding to the DHCP Binding Database

A switch maintains a DHCP binding database, which is used for dynamic IP lockdown as well as for DHCP and ARP packet validation. The DHCP snooping feature maintains the lease database by learning the IP-to-MAC bindings of VLAN traffic on untrusted ports. Each binding consists of the client MAC address, port number, VLAN identifier, leased IP address, and lease time.

Dynamic IP lockdown supports a total of 4K static and dynamic bindings with up to 64 bindings per port. When DHCP snooping is enabled globally on a VLAN, dynamic bindings are learned when a client on the VLAN obtains an IP address from a DHCP server. Static bindings are created manually with the CLI or from a downloaded configuration file.

When dynamic IP lockdown is enabled globally or on ports the bindings associated with the ports are written to hardware. This occurs during these events:

- Switch initialization
- Hot swap
- A dynamic IP lockdown-enabled port is moved to a DHCP snooping-enabled VLAN
- DHCP snooping or dynamic IP lockdown characteristics are changed such that dynamic IP lockdown is enabled on the ports

Potential Issues with Bindings

- When dynamic IP lockdown enabled, and a port or switch has the maximum number of bindings configured, the client DHCP request will be dropped and the client will not receive an IP address through DHCP.
- When dynamic IP lockdown is enabled and a port is configured with the maximum number of bindings, adding a static binding to the port will fail.
- When dynamic IP lockdown is enabled globally, the bindings for each port are written to hardware. If global dynamic IP lockdown is enabled and disabled several times, it is possible to run out of buffer space for additional bindings. The software will delay adding the bindings to hardware until resources are available.

Adding a Static Binding

To add the static configuration of an IP-to-MAC binding for a port to the lease database, enter the **ip source-binding** command at the global configuration level. Use the **no** form of the command to remove the IP-to-MAC binding from the database.

Syntax: [no] ip source-binding <vlan-id> <ip-address> <mac-address> <port-number>

vlan-id Specifies a valid VLAN ID number to bind with the specified MAC and IP addresses on the port in the DHCP binding database.

ip-address Specifies a valid client IP address to bind with a VLAN and MAC address on the port in the DHCP binding database.

mac-address Specifies a valid client MAC address to bind with a VLAN and IP address on the port in the DHCP binding database.

port-number Specifies the port number on which the IP-to-MAC address and VLAN binding is configured in the DHCP binding database.

Note

Note that the **ip source-binding** command is the same command used by the Dynamic ARP Protection feature to configure static bindings. The Dynamic ARP Protection and Dynamic IP Lockdown features share a common list of source IP-to-MAC address bindings.

Verifying the Dynamic IP Lockdown Configuration

To display the ports on which dynamic IP lockdown is configured, enter the **show ip source-lockdown status** command at the global configuration level.

Syntax: show ip source-lockdown status

An example of the **show ip source-lockdown status** command output is shown in Figure 20. Note that the operational status of all switch ports is displayed. This information indicates whether or not dynamic IP lockdown is supported on a port.

```
ProCurve(config)# show ip source-lockdown status
Dynamic IP Lockdown (DIPLD) Information

Global State: Enabled

      Port      Operational State
-----  -
A1        Active
A2        Not in DHCP Snooping vlan
A3        Disabled
A4        Disabled
A5        Trusted port, Not in DHCP Snooping vlan
. . . . .
```

Figure 20. Example of show ip source-lockdown status Command Output

Displaying the Static Configuration of IP-to-MAC Bindings

To display the static configurations of IP-to-MAC bindings stored in the DHCP lease database, enter the **show ip source-lockdown bindings** command.

Syntax: show ip source-lockdown bindings [*<port-number>*]
port-number (Optional) Specifies the port number on which source IP-to-MAC address and VLAN bindings are configured in the DHCP lease database.

An example of the **show ip source-lockdown bindings** command output is shown in [Figure 21](#).

```
ProCurve(config)# show ip source-lockdown bindings

Dynamic IP Lockdown (DIPLD) Bindings

Mac Address      IP Address      VLAN    Port    Not in HW
-----
001122-334455   10.10.10.1     1111    X11
005544-332211   10.10.10.2     2222    Trk11   YES
. . . . .
```

Figure 21. Example of show ip source-lockdown bindings Command Output

In the **show ip source-lockdown bindings** command output, the “Not in HW” column specifies whether or not (YES or NO) a statically configured IP-to-MAC and VLAN binding on a specified port has been combined in the lease database maintained by the DHCP Snooping feature.

Debugging Dynamic IP Lockdown

To enable the debugging of packets dropped by dynamic IP lockdown, enter the **debug dynamic-ip-lockdown** command.

Syntax: debug dynamic-ip-lockdown

To send command output to the active CLI session, enter the **debug destination session** command.

Counters for denied packets are displayed in the **debug dynamic-ip-lockdown** command output. Packet counts are updated every five minutes. An example of the command output is shown in [Figure 22](#).

When dynamic IP lockdown drops IP packets in VLAN traffic that do not contain a known source IP-to-MAC address binding for the port on which the packets are received, a message is entered in the event log.


```
ProCurve(config)# debug dynamic-ip-lockdown

DIPLD 01/01/90 00:01:25 : denied ip 192.168.2.100 (0)
(PORT 4) -> 192.168.2.1 (0), 1 packets
DIPLD 01/01/90 00:06:25 : denied ip 192.168.2.100 (0)
(PORT 4) -> 192.168.2.1 (0), 294 packets
DIPLD 01/01/90 00:11:25 : denied ip 192.168.2.100 (0)
(PORT 4) -> 192.168.2.1 (0), 300 packets
DIPLD 01/01/90 00:16:25 : denied ip 192.168.2.100 (0)
(PORT 4) -> 192.168.2.1 (0), 300 packets
DIPLD 01/01/90 00:21:25 : denied ip 192.168.2.100 (0)
(PORT 4) -> 192.168.2.1 (0), 299 packets
DIPLD 01/01/90 00:26:25 : denied ip 192.168.2.100 (0)
(PORT 4) -> 192.168.2.1 (0), 300 packets
DIPLD 01/01/90 00:31:25 : denied ip 192.168.2.100 (0)
(PORT 4) -> 192.168.2.1 (0), 300 packets
DIPLD 01/01/90 00:36:25 : denied ip 192.168.2.100 (0)
(PORT 4) -> 192.168.2.1 (0), 299 packets
DIPLD 01/01/90 00:41:25 : denied ip 192.168.2.100 (0)
(PORT 4) -> 192.168.2.1 (0), 300 packets
DIPLD 01/01/90 00:46:25 : denied ip 192.168.2.100 (0)
(PORT 4) -> 192.168.2.1 (0), 300 packets
DIPLD 01/01/90 00:51:25 : denied ip 192.168.2.100 (0)
(PORT 4) -> 192.168.2.1 (0), 300 packets
DIPLD 01/01/90 00:56:25 : denied ip 192.168.2.100 (0)
(PORT 4) -> 192.168.2.1 (0), 299 packets
DIPLD 01/01/90 01:01:25 : denied ip 192.168.2.100 (0)
(PORT 4) -> 192.168.2.1 (0), 300 packets
```

Figure 22. Example of debug dynamic-ip-lockdown Command Output

Release M.10.44 through M.10.64 Enhancements

Software fixes only, no new enhancements.

Release M.10.65 Enhancements

Release M.10.65 includes the following enhancement:

- **Enhancement (PR_0000001316)** — The MSTP VLAN Assignment is enhanced.

MSTP VLAN Configuration Enhancement

Caution

When this software version is installed, the prior VLAN ID-to-MSTI mappings do not change. However, this enhancement is not backward-compatible. If you install a software version prior to this version, and you have configured MSTI entries instances mapped to VLANs, they will be removed from the configuration file when booting to the prior version of software. You must remove all MSTP mappings from the config file and then reconfigure the instance mapping after you are running the desired software version.

If you want to save a copy of the switch configuration file to a tftp server before you install this version, use this command:

```
ProCurve(config)# copy startup-config tftp <ip-address tftp server>  
                    <filename>
```

where <filename> is the name you choose for the pre-enhancement configuration file that you are saving.

Overview

The MSTP VLAN configuration enhancement allows you to preconfigure an MSTP regional topology and ensure that the same VLAN ID-to-MSTI assignments exist on each MSTP switch in the region.

The default behavior of the **spanning-tree instance vlan** command changes so that, before a static VLAN is configured or a dynamic VLAN is learned on the switch, you can preconfigure its VLAN ID-to-MSTI mapping. Later, when the VLAN is created, it is automatically assigned to the MSTI to which you had previously mapped it.

By supporting preconfigured VLAN ID-to-MSTI topologies, the MSTP VLAN configuration enhancement provides the following benefits:

- **Scalability:** In a network design in which you plan to use a large number of VLANs, you can preconfigure identical VLAN ID-to-MSTI mappings on all switches in a single, campus-wide MST region, regardless of the specific VLANs that you later configure on each switch. After the initial VLAN ID-to-MSTI mapping, you can decide on the exact VLANs that you need on each switch.

All switches in a region must be configured with the same VLAN ID-to-MSTI mappings and the same MSTP configuration identifiers (region name and revision number).

- **Flexibility:** By preconfiguring identical VLAN ID-to-MSTI mappings on all switches in an MST region, you can combine switches that support different maximum numbers of VLANs.
- **Network stability:** You can reduce the interruptions in network connectivity caused by the regeneration of spanning trees in the entire network each time a configuration change in VLAN-to-MSTI mapping is detected on a switch. The negative impact on network performance is reduced if all newly created VLANs are pre-mapped to the correct MST instances. Later, VLAN creation and deletion are ignored by MSTP and no interruption in spanning-tree traffic occurs.
- **Usability:** Dynamically learned GVRP VLANs can be mapped to MSTIs and support MSTP load balancing.

Enabling MSTP on the Switch

If you have not enabled MSTP on the switch, you must enable it to use this feature. To enable MSTP, perform these steps.

1. Enter the command to enable MSTP:

```
ProCurve(config)# spanning-tree protocol-version mstp
```

You will see this message:

```
STP version was changed. To activate the change you must save the
configuration to flash and reboot the device.
```

2. Save the configuration change to flash.

```
ProCurve(config)# write mem
```

3. Reboot the switch.

```
ProCurve(config)# reload
```

```
Device will be rebooted, do you want to continue [y/n]? y
```

PreConfiguring VLANs in an MST Instance

When you configure an MSTP regional topology, you create multiple spanning-tree instances. Each MST instance provides a fully connected active topology for a particular set of VLANs.

Each switch in an MSTP region is configured with the following set of common parameters:

- Region name (**spanning-tree config-name**)
- Region revision number (**spanning-tree config-revision**)
- Identical VLAN ID-to-MSTI mapping (**spanning-tree instance vlan**)

Each MST instance supports a different set of VLANs. A VLAN that is mapped to an MST instance cannot be a member of another MST instance.

The MSTP VLAN Configuration enhancement allows you to ensure that the same VLAN ID-to-MSTI assignments exist on each MSTP switch in a region. Before a static VLAN is configured or a dynamic VLAN is learned on the switch, you can use the **spanning-tree instance vlan** command to map VLANs to each MST instance in the region. Later, when the VLAN is created, the switch automatically assigns it to the MST instance to which you had previously mapped it.

Syntax: [no] spanning-tree instance < 1..16 > vlan < vid [vid..vid] >
no spanning-tree instance < 1..16 >

Configuring MSTP on the switch automatically configures the IST instance and places all statically and dynamically configured VLANs on the switch into the IST instance. This command creates a new MST instance (MSTI) and moves the VLANs you specify from the IST to the MSTI.

*You must map at least one VLAN to an MSTI when you create it. You cannot map a VLAN ID to more than one instance. You can create up to 16 MSTIs in a region. The **no** form of the command removes one or more VLANs from the specified MSTI. If no VLANs are specified, the **no** form of the command deletes the specified MSTI. When you remove a VLAN from an MSTI, the VLAN returns to the IST instance, where it can remain or be re-assigned to another MSTI configured in the region.*

Note: *The valid VLAN IDs that you can map to a specified MSTI are from 1 to 4094. The VLAN ID-to-MSTI mapping does not require a VLAN to be already configured on the switch. The MSTP VLAN enhancement allows you to preconfigure MSTP topologies before the VLAN IDs associated with each instance exist on a switch.*

When you use preconfigured VLAN ID-to-MSTI topologies, ensure that MSTP switches remain in the same region by mapping all VLAN IDs used in the region to the same MSTIs on each regional switch.

Configuring MSTP Instances with the VLAN Range Option

If you use the **spanning-tree instance** command with the VLAN range option, even if the range includes VLANs that are not currently present on the switch, the entire range of VLANs is configured. For example, if VLANs 1, 5, and 7 are currently present and you enter this command:

```
ProCurve(config)# spanning-tree instance 1 vlan 1-10
```

then all the VLANs from 1 through 10 are included, even those VLANs that are not present.

[Figure 23](#) shows an example of an MSTP instance configured with the VLAN range option. All the VLANs are included in the instance whether they exist or not.

```
ProCurve(config)# show spanning-tree mst-config

MST Configuration Identifier Information
MST Configuration Name: MSTP1
MST Configuration Revision: 1
MST Configuration Digest: 0x51B7EBA6BEED8702D2BA4497D4367517

IST Mapped VLANs :

Instance ID Mapped VLANs
-----
1           1-10
```

Figure 23. Example of Mapping VLANs with the Range Option where all VLANs are Included

Note

If you want all switches to be in the same MST region, they should all have a software version that supports this enhancement installed, or have the same VLANS configured.

It is likely that switches with a VLAN range configured prior to this enhancement and switches with a VLAN range configured after updating the switch with this enhancement will have different Configuration Digests.

The Common Spanning Tree (CST) will still have the correct root associations.

Operating Notes

- Configuring MSTP on the switch automatically configures the Internal Spanning Tree (IST) instance and places all statically and dynamically configured VLANs on the switch into the IST instance. The **spanning-tree instance vlan** command creates a new MST instance and moves the VLANs you specify from the IST to the MSTI.

You must map a least one VLAN ID to an MSTI when you create it. You cannot map a VLAN ID to more than one instance. You can create up to 16 MSTIs in a region.

- The **no** form of the **spanning-tree instance vlan** command removes one or more VLANs from the specified MSTI. If no VLANs are specified, the **no** form of the command deletes the specified MSTI.

When you remove a VLAN from and MSTI, the VLAN returns to the IST instance, where it can remain or be re-assigned to another MSTI configured in the region.

Enhancements

Release M.10.66 Enhancements

- If you enter the **spanning-tree instance vlan** command before a static or dynamic VLAN is configured on the switch to preconfigure VLAN ID-to-MSTI mappings, no error message is displayed. Later, each newly configured VLAN that has already been associated with an MSTI is automatically assigned to the MSTI.

This new default behavior differs from automatically including configured (static and dynamic) VLANs in the IST instance and requiring you to manually assign individual static VLANs to an MSTI.

- The valid VLAN IDs that you can map to a specified MSTI are from 1 to 4094. The VLAN ID-to-MSTI mapping does not require a VLAN to be already configured on the switch. The MSTP VLAN enhancement allows you to preconfigure MSTP topologies before the VLAN IDs associated with each instance exist on a switch.
- When you use preconfigured VLAN ID-to-MSTI topologies, ensure that MSTP switches remain in the same region by mapping all VLAN IDs used in the region to the same MSTIs on each regional switch.
- When you update switch software, the existing MSTP topology configuration is automatically saved. All existing VLAN ID-to-MSTI assignments are maintained on a switch for uninterrupted MSTP network operation.

Release M.10.66 Enhancements

Release M.10.66 includes the following enhancement:

- **Enhancement (0000000818)** — This enhancement allows syslog configuration via SNMP.

Configure Logging via SNMP

Debug messages generated by the software can be sent to a syslog server. This feature provides the ability to enter addresses and filter parameters for syslog using SNMP, which allows more options for remote access and management of the switch. The HP enterprise MIB `hpicfSyslog.mib` is added to allow the configuration and monitoring of syslog. (RFC 3164 supported)

The CLI has some additional parameters that permit interoperability with SNMP that are explained below.

Note

See the section [“Command Differences for the ProCurve Series 2600/2800/3400cl/6400cl Switches”](#) on page 142 for command differences on these switches.

Adding a Description for a Syslog Server

You can associate a user-friendly description with each of the IP addresses (IPv4 only) configured for syslog using the CLI or SNMP. The CLI command is:

Syntax: logging <ip-addr> control-descr <text_string>
no logging <ip-addr> [control-descr]

*An optional user-friendly description that can be associated with a server IP address. If no description is entered, this is blank. If <text_string> contains white space, use quotes around the string. IPv4 addresses only. Use the **no** form of the command to remove the description.*

Limit: 255 characters

Note: To remove the description using SNMP, set the description to an empty string.

```
ProCurve(config)# logging 10.10.10.2 control-descr syslog_one
```

Figure 29. Example of the Logging Command with a Control Description

Caution

Entering the **no logging** command removes ALL the syslog server addresses without a verification prompt.

Adding a Priority Description

You can add a user-friendly description for the set of syslog filter parameters using the **priority-descr** option. The description can be added with the CLI or SNMP. The CLI command is:

Syntax: logging priority-descr <text_string>
no logging priority-descr

*Provides a user-friendly description for the combined filter values of **severity** and **system module**. If no description is entered, this is blank. If <text_string> contains white space, use quotes around the string. Use the **no** form of the command to remove the description.*

Limit: 255 characters

```
ProCurve(config)# logging priority-descr severe-pri
```

Figure 30. Example of the Logging Command with a Priority Description

Note

A notification is sent to the SNMP agent if there are any changes to the syslog parameters either through the CLI or with SNMP.

Command Differences for the ProCurve Series 2600/2800/3400cl/6400cl Switches

CLI Commands. The ProCurve series 2600/2800/3400cl/6400cl switches do not have the following CLI logging commands:

- **logging severity**
- **logging system-module**

SNMP Commands. The ProCurve series 2600/2800/3400cl/6400cl switches do not support the following SNMP objects:

- hpicfSyslogPrioritySeverity
- hpicfSyslogSystemModule

Operating Notes

- Duplicate IP addresses are not stored in the list of syslog servers.
- If the default severity value is in effect, all messages that have severities greater than the default value are passed to syslog. For example, if the default severity is “debug”, all messages that have severities greater than debug are passed to syslog.
- There is a limit of six syslog servers. All syslog servers are sent the same messages using the same filter parameters.
- An error is generated for an attempt to add more than six syslog servers.

Release M.10.67 Enhancements

Software fixes only, no new enhancements.

Release M.10.68 Enhancements

Release M.10.68 includes the following enhancement:

- **Enhancement (PR_0000003127)** — A Link Trap and LACP Global enable/disable feature has been added.

LACP and Link Traps Global Disable

Two SNMP commands are added to allow disabling of LACP and link traps on multiple ports at one time. The new commands operate in the same manner as the CLI commands **no int all lacp** and **no snmp-server enable traps link-change all**.

The new SNMP OIDs are:

```
hpSwitchLACPConfig OBJECT IDENTIFIER ::= { hpSwitchConfig 28 }
```

```
hpSwitchLACPAllPortsStatus OBJECT-TYPE
```

```
    SYNTAX INTEGER {  
  
        disabled (1),  
        active (2),  
        passive (3)  
  
    }
```

```
    ACCESS read-write
```

```
    STATUS mandatory
```

```
    DESCRIPTION "Used to set administrative status of LACP on all the  
        ports. A Port can have one of the three  
        administrative status of LACP.  
        Active/Passive/Disabled are the three states."
```

```
 ::= { hpSwitchLACPConfig 1 }
```

Enhancements

Release M.10.69 Enhancements

```
hpSwitchLinkUpDownTrapAllPortsStatus OBJECT-TYPE
    SYNTAX INTEGER {
        enable (1),
        disable (2)
    }
    ACCESS read-write
    STATUS current
    DESCRIPTION "Used to either enable/disable the Link Up/Link Down traps
        for all the ports."
    ::= { hpSwitchPortConfig 3 }
```

Release M.10.69 Enhancements

Release M.10.69 includes the following enhancement (Not a public release).

- **Enhancement (PR_0000010783b)** — The CLI output for the **show tech transceivers** command has been enhanced to be consistent with other platforms.

Release M.10.70 Enhancements

Software fixes only, no new enhancements.

Release M.10.71 Enhancements

Release M.10.71 includes the following enhancement (Not a public release).

- **Enhancement (PR_0000011636)** — This enhancement adds the client's IP address to the RADIUS accounting packets sent to the RADIUS server by the switch. The IP address of the client is included in the RADIUS accounting packet sent by the switch to the RADIUS server. The client obtains the IP address through DHCP, so DHCP snooping must be enabled for the VLAN of which the client is a member.

Release M.10.72 Enhancements

Software fixes only, no new enhancements.

Software Fixes in Release M.08.51 - M.10.72

Software fixes are listed in chronological order, oldest to newest. To review the list of fixes included since the last general release that was published, go to [“Release M.10.21” on page 165](#).

Unless otherwise noted, each new release includes the software fixes added in all previous releases. Release M.08.51 was the first software release for the HP ProCurve 3400cl Series.

Release M.08.52

Updated Boot ROM image to I.08.02 to address Manufacturing test condition.

Release M.08.53 (Never Released)

- Updated Boot ROM image to I.08.03 to address Manufacturing test condition.

Release M.08.54

- First release to fully support LR and CX4 transceivers installed in the optional cl Module (J8434A and J8435A).

Release M.08.55 - Release M.08.60

Releases M.08.55 through M.08.60 were never built.

Release M.08.61

Problems Resolved in Release M.08.61

- **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.
- **ACL (PR_1000207620)** — TCP and UDP traffic is sometimes incorrectly permitted through an ACL.
- **CDP (PR_1000195343)** — Entering the command **show cdp neighbor detail x** (where **x** is the port number) displays details for all active ports with CDP neighbors whose numbers begin with **x**. Only occurs when the **detail** parameter is included.
- **CDP/LLDP (PR_1000201275)** — The CDP/LLDP text output has been improved with the following two changes:
 1. In **show LLDP config** output, LLDP refresh interval is renamed to LLDP Transmit Interval.

2. In **show CDP** the Yes is changed to Yes, (Receive Only).
- **CLI (PR_1000192677)** — **Show access-list ports <tab>** does not list the all keyword. The command only shows [PORT-LIST] as input for the command.
 - **Console/TELNET (PR_1000195647)** — When a console or TELNET session hangs, issuing the 'kill' command also hangs.
 - **Crash (PR_1000193582)** — Software Exception when clicking on the Identity Tab of a Member Switch in the Web user interface. The switch may crash with a message similar to:
Software exception at http_state.c:1138 in 'mHttpCtrl' TaskID = 0x1722cf8.
 - **Crash (PR_1000196129)** — Removing the J8434A module (10Gigabit) creates configuration problems and sometimes causes the switch to crash.
 - **Crash (PR_1000199535)** — Sometimes the command **show boot-history** results in a bus error. (pre-release)
 - **Crash (PR_1000201614)** — When the switch is set with a 16-character manager password in the setup menu, a 'Bus error' crash may occur. The bus errors vary.
 - **Crash (PR_1000204782)** — Bus error when copying a configuration to the switch. The switch may crash with a message similar to:
Bus error: HW Addr=0x594f5531 IP=0x004ff8a8 Task='mftTask' Task ID=0x126eba0 fp: 0x00000000 sp:0x0126e7d0 lr:0x001e655c.
 - **IP Addmgr (PR_1000200338)** — CPU-based protocol stops working. The memory corruption of text caused many tasks to hang or be SUSPENDED, since the switch is trying to execute invalid instructions.
 - **MIB (PR_1000206519)** — The RFC 3636 MIB implemented is not correct. (pre-release)
 - **Open VLAN (PR_1000210932)** — Open VLAN mode (Unauth VLAN) does not work with any Port-Security Learn-Mode.
 - **Port toggle (PR_1000216940)** — 10 Gigabit, CX4 port toggles (that is, Link up, Link down, and so on). (pre-release)
 - **QOS (PR_1000200746)** — Configure a dscp-map name that requires quotes such as "Code Point 0". Save this name in the configuration file and reboot the switch, the name is truncated to "Code".
 - **QOS (PR_1000213489)** — The command **show QOS resources** displays blank information for the 10 Gigabit module.
 - **SNMP (PR_1000196170)** — Traps are not buffered before the IP stack is initialized, causing the possibility of missing some traps generated during startup.
 - **Syslog (PR_1000215699)** — Pre-boot event log messages are cached for syslog and syslog is only able to send those cached entries after the switch fully boots. The cache size is limited, so in some cases, not all event log messages will be sent via syslog.
-

- **Web UI (PR_1000177915)** — Device View from the Web user interface is missing.
- **Web UI/Port Security (PR_1000195894)** — The Web user interface does not allow the user to select multiple ports when configuring port-security.

Release M.08.62

Problems Resolved in Release M.08.62

- **Crash (PR_1000207542)** — The switch may crash with a bus error or task hang.
- **Crash (PR_1000216170)** — The switch crashes with an `mftTask Bus Error` whenever a user attempts to upload the startup-configuration from a TFTP server. The switch accepts the command with no errors, however the system immediately crashes after the reboot.
- **Jumbo/Flow control (PR_1000217576)** — When the switch is configured for both flow control and jumbo packets, an Error Message is not generated as stated in the instruction manual.
- **Port Security (PR_1000203984)** — When the limit is reached, the warning message is displayed: `Number of configured addresses on port xx exceeds address-limit`. The address is saved and displayed in the address list of **Show Port-security xx**. Data from the added address is passed by the switch.

Release M.08.63

Problems Resolved in Release M.08.63 (Not a general release)

- **Crash (PR_1000205768)** — A `null` System Name in the Web user interface may crash the switch with a message similar to:

```
"Software exception at lldpSysNameTlv.c:251 -- in 'mllldpCtrl', task ID = 0x12dc88 -> ASSERT: failed".
```
- **Web UI (PR_93721)** — The web user interface Status screen does not display all ports, and the scroll bar does not work.
- **Web UI (PR_1000191635)** — The Port column may not be sorted correctly in all Web user interface screens.
- **XRRP (PR_1000217651)** — Running different XRRP versions causes excessive event log messages like:

```
Rcvd a pkt with version number 2, expected 1  
Remote rtr 2 domain 2 is miss-configured.
```
- **Crash (PR_1000217354)** — Bus error in `mSnmppCtrl` task when adding a less-specific route and adding it again through the CLI.

Release M.08.64

Problems Resolved in Release M.08.64 (Not a general release)

- **IP Routing (PR_1000220668)**— Fatal exception when routing with more than 8 trunks configured and IP routing enabled.

Release M.08.65

Problems Resolved in Release M.08.65 (Never released)

- **Crash (PR_1000194486)** — The switch may crash with a message similar to:
`Software exception at bcm 1 CpuLearn.c:1308.`
- **Counters (PR_1000221089)** — The 64 bit counters may not always be correct.
- **Counters (PR_1000219548)** — Collision counters do not increment accurately.

Release M.08.66

Problems Resolved in Release M.08.66 (Not a general release)

- **PPMGR (PR_1000225645)**— The ProCurve 10GbE X2-SC SR Optic (J8436A) transceiver fails self test on boot up when installed in slot B/8.

Release M.08.67

Problems Resolved in Release M.08.67 (Not a general release)

- **Authentication (PR_1000217338)** — Inconsistent authentication results with EAP-TLS and EAP-PEAP authorization types.
- **Config (PR_1000207697)** — Loading a startup-configuration file fails when attempting to declare a VLAN in the configuration file as a management VLAN, and the VLAN does not currently exist on the switch. The switch indicates the downloaded file as being corrupted, listing the `vid` of the specified management VLAN as not being found.
- **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.
- **Web UI (PR_1000214188)** — The scroll bar does not display or respond correctly after resizing a window.

Release M.08.68

Problems Resolved in Release M.08.68 (Not a general release)

- **Switching (PR_1000232312)** — In cases where traffic is being L2 switched or L3 routed from one port at Gigabit speeds to a group of ports (i.e. to a VLAN) where one of the outbound ports is running at a slower speed, traffic may have been dropped even to egress ports running at Gigabit speeds. This PR addresses the dropped packets for the Gig-to-Gig port traffic. Gig-to-100Mbps transfers may still experience packet drops due to congestion (as is normal in any oversubscribed scenario).

Release M.08.69

Problems Resolved in Release M.08.69

- **802.1s STP (PR_1000229407)** - Edge port configuration is lost after the configuration file is transferred using TFTP.
- **802.1X (PR_1000208530)** - Switch may crash with 802.1X configured, with a message similar to:

```
Crash: aaa8021x_init dereferencing a null pointer, writing to low memory
```
- **CLI (PR1000202435)** — “show config” does not show IGMP fast-leave configuration.
- **Config (PR_94943)** — Setup Screen allows Proxy-ARP configuration when IP routing is disabled
- **Config (PR1000216051)** — Copying a previously saved startup-configuration with “stack join (mac address)” to a member switch of the IP stack will break the membership of that stack.
- **Crash (PR_1000229656)** - switch crashes when RADIUS is unavailable.
- **Crash (PR_1000233993)** - Switch may crash with a message similar to:

```
Software exception at exception.c:373 -- in 'mSnmppCtrl', task ID = 0x5b85fd0 -> Memory system error.
```
- **Crash (PR_1000239085)** - The switch may crash with a message similar to:

```
Software exception at esi_stacking.c:2578 -- in 'tHttpd'.
```
- **DHCP (PR1000207419)** — DHCP Relay agent is disabled by default.
- **IP Helper/DHCP Relay (PR_1000197046)** - IP helper may not handle "DHCP Inform" relay properly.
- **Menu (PR_1000221018)** - Setup Menu allows Proxy-ARP configuration when IP routing is disabled.

Software Fixes in Release M.08.51 - M.10.72

Release M.08.70

- **Port Security (PR_1000203984)** — CLI port-security "mac-address" command will save address above the limit.
- **SNMP (PR_1000212170)** — The Switch transmits Warm and Cold Start traps with an agent address of 0.0.0.0.
- **Spanning Tree (PR_1000214598)** - The switch will not accept the spanning-tree 1 mode fast command within the CLI.
- **System Hang (PR_1000200341)** - Added an exception handler to prevent a case where the system may hang.
- **XRRP (PR_1000217922)** — XRRP router may fail back to the XRRP peer router even with Infinite Failback enabled.

Release M.08.70

Problems Resolved in Release M.08.70 (Not a general release)

- **ACL (PR_1000213663)** — When configuring ACLs, the Switch incorrectly reports:
Duplicate access control entry.
- **Broadcast throttling (PR_1000240494)** — Broadcast throttling does not work correctly on Gigabit/second and 10-Gigabit/second ports.
- **Mesh (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.
- **MIB (PR_1000236875)** — The switch is reporting etherType/size errors as part of "ifInDiscards," but the packets are not really dropped.
- **Packet buffers (PR_1000237366)** — Improved packet buffer allocation for better data handling.
- **Self-test (PR_1000239302)** — The Switch reports a false self-test failure when a J8436A SR transceiver is installed in Port B of a J8435A 10-GbE Media Flex module.
- **Web/Stack Mgmt (PR_1000239924)** — As an IP Stack Management Commander, the Switch does not display the device view (back of box) for a 2626 switch that is a member.

Release M.08.71

Problems Resolved in Release M.08.71 (Never released)

- **Crash (PR_1000232283)** — The switch may crash with a message similar to:
Software exception at fileTransferTFTP.c:182 -- in 'mftTask', task ID = 0x107ee0.

- **LLDP (PR_1000241315)** — CLI command "show LLDP" does not display information correctly.
- **Web Auth (PR_1000230444)** — Using port-based web authentication on the Switch will cause some users to never receive the web authentication screen. This occurs if a client receives the same unauthenticated DHCP address that a previous authorized client has used.
- **802.1s (PR_1000233920)** — 802.1s (MSTP) blocks a port that is connected to an RSTP device.

Release M.08.72

Problems Resolved in Release M.08.72 (Not a general release)

- **Crash (PR_1000234773)** — The switch may crash with a message similar to:

```
"ifInfo" task: SubSystem 0 went down: 01/01/90 00:03:16 NMI event  
SW:IP=0x004c1bdc MSR:0x0000b032 LR:0x004c3850 Task='ifInfo' Task  
ID=0x137c980 cr: 0x22242040 sp:0x0137bef8 xer:0x00000000.
```
- **Flow Control (PR_1000241296)** — Switch was unable to support flow control between any ingress and any egress ports.
- **SNMP (PR_1000003378)** — SNMP switch time may drift with event log updates occurring every 1.5 hours.
- **Web UI (PR_1000211978)** — On a Stack Management Commander, when using "stack access" to view members, the screen does not display correct information.

Release M.08.73

Problems Resolved in Release M.08.73 (Not a general release)

- **Crash (PR_1000282197)** — The 3400cl-48G may experience crash or reboot symptoms on initial install of the switch. The crashes have a PPC crash heading. The switch may reboot with no crash history, simply the following message:

System reboot due to power failure.

Boot ROM — Updated to I.08.07 version to support fix for PR 1000282197.

Release M.08.74

Problems Resolved in Release M.08.74 (Not a general release)

- **Meshing (PR_1000282427)** — Multicast traffic not forwarded out 10 Gigabit mesh ports.

Release M.08.75

Problems Resolved in Release M.08.75

- **LR optic (PR_1000282195)** — After a switch reboot, certain 10GbE X2-SC LR Optic (J8437A) transceivers will lose its configuration. Administrator will be unable to turn off LACP, and CLI commands will not be displayed.
- **XRRP (PR_1000280213)** — When configuring a XRRP instance, although the subnet is configured properly, the following error message is logged:

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

Release M.08.76

Problems Resolved in Release M.08.76 (Never released)

- **IP Routing (PR_1000254254)** — L3 address table is not learned correctly from unsolicited ARPs.
- **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 the switch.
- **TCP (PR_1000246186)** — Switch is susceptible to VU#498440.
- **VLAN (PR_1000214406)** — When trying to delete a VLAN created as a management VLAN, the switch fails to remove the management VLAN statement from the running configuration file.

Web UI (PR_1000284653) — When using the web user interface "IP Stack Management", and there are more than 100 potential Members present on a VLAN, the Switch will learn new potential Members, but deletes previously learned Members.

Release M.08.77

Problems Resolved in Release M.08.77 (Not a general release)

- **ACL (PR_1000283338)** — The commands "show port-access mac" and "show port-access web" incorrectly display the number of clients authenticated.
- **Meshing (PR_1000219337)** — Unstable RSTP topology when root switch is power-cycled and connected to a mesh.

Release M.08.78

Problems Resolved in Release M.08.78 (Not a general release)

- **Enhancement (PR_1000291806)** — Fast boot enhancement.
- **MSTP (PR_1000286883)** — Slow MSTP fail-over and fall-back time.

Release M.08.79

Problems Resolved in Release M.08.79 (Not a general release)

- **Fault (PR_1000089786)** — Chassis fault LED stops blinking after a new OS image was downloaded to the switch.
- **Ports (PR_1000090867)** — The dual personality ports (RJ-45 and mini-GBIC) lose state (running speed) after being hot swapped in or out.
- **Enhancement (PR_100292455)** — Rate display for ports on CLI. New command: "show interface port-utilization", not available on Menu nor Web Interface.

Release M.08.80

Problems Resolved in Release M.08.80 (Never released)

- **RSTP (PR_1000297195)** — The switch repeatedly flushes its MAC address table, resulting in intermittent flooding of all traffic.

Release M.08.81

Problems Resolved in Release M.08.81 (Not a general release)

- **XRRP (PR_1000291250)** — When a XRRP router is rebooted and activates its virtual MAC address, it incorrectly transmits ARP requests, which fails to update forwarding tables and ARP caches.

Release M.08.82

Problems Resolved in Release M.08.82 (Not a general release)

- **Meshing (PR_1000300165)** — Packets larger than 1482 bytes within a mesh will be reported as FCS receive errors and may generate excessive CRC error messages in the event log.

- **RSTP (PR_1000300623)** — Under some circumstances, the switch may allow packets to loop for an extended period of time.

Release M.08.83

Problems Resolved in Release M.08.83 (Not a general release)

- **Crash (PR_1000297510)** — When using the Web User Interface and the switch is set as commander for stacking, the switch may crash.
- **Event Log/ARP (PR_1000293466)** — Generic Link Up message not showing up and unnecessary flushing of ARP cache.
- **KMS (PR_1000287934)** — Some Key Management System (KMS) configuration commands have no effect.
- **Setup (PR_1000301498)** — Manual IP address can not be set using "setup" menu. (pre-release)

Release M.08.84

Problems Resolved in Release M.08.84 (Never released)

- **CLI Enhancement (PR_1000306695)** — Added "show tech transceivers" to display Serial Number information for installed mGBIC and 10Gig X2 transceivers. Allows removable transceiver serial numbers to be read without removal of the transceivers from the switch.

Release M.08.85

Problems Resolved in Release M.08.85 (Never released)

- **RSTP (No PR)** — Resolved broadcast storm caused by an unstable RSTP topology.

Release M.08.86

Problems Resolved in Release M.08.86

- **CLI/DHCP (PR_1000286898)** — Under some conditions, the CLI may freeze or lock up.
- **IGMP (PR_1000301557)** — Data-driven IGMP does not prevent flooding when no IP address exists on a VLAN.
- **RSTP (PR_1000306227)** — RSTP TCNs cause high CPU utilization and slow software based routing.

- **SNMP (PR_1000295753)** — Removing 'public' SNMP community generates an empty Event Log message.

Release M.08.87

Problems Resolved in Release M.08.87 (Not a general release)

- **Crash/STP (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 'mSnmprCtrl', task ID =  
0x12d14b8\n-> ASSERT: failed.
```
- **Menu (PR_1000306213)** — When using the Menu to create a trunk, the new trunk ports will become disabled after a switch reboot.
- **OSPF (PR_1000280427)** — OSPF MD5 Authentication failure.
- **RSTP (PR_1000309683)** — Temporary routing or switching problems after RSTP is disabled.

Release M.08.88

Problems Resolved in Release M.08.88 (Not a general release)

- **CLI (PR_1000310849)** — Under a heavy load where packets received on a 10-Gigabit port are dropped, the RX drop counter values decrease when they should increase.
- **LLDP (PR_1000310666)** — The command "show LLDP" does not display information learned from CDPv2 packets.
- **SNMP Traps (PR_1000285195)** — Switch does not save the option to disable a Link up/down SNMP trap after a switch reboot.
- **Web /Stacking (PR_1000308933)** — Added Web User Interface stacking support for the new Series 3500y1 switches, providing a 3500y1 "back-of-box" display when the 3400cl or 6400cl is stack commander and a 3500y1 is a stack member.

Release M.08.89

Problems Resolved in Release M.08.89 (Never released)

- **Enhancements (PR_1000313819)** — Added two enhancements:
 - DNS Names for Ping and Traceroute

Software Fixes in Release M.08.51 - M.10.72
Release M.08.90

- **RADIUS Configuration via SNMP.** For details refer to [“Using SNMP To View and Configure Switch Authentication Features”](#) on page 35.
- **Port Security (PR_1000304202)** — The port-security MAC address learn mode does not function correctly between 'port-security' ports.
- **SNMP (PR_1000310841)** — User can assign illegal values for CosDSCPpolicy through SNMP. All other user-interfaces for configuring QoS function correctly.

Release M.08.90

Problems Resolved in Release M.08.90 (Not a general release)

- **Crash/log (PR_1000282359)** - When searching the log for an extremely long string, the switch may crash with a bus error similar to:

```
PPC Bus Error exception vector 0x300: Stack Frame=0x0c8c1a70 HW  
Addr=0x6a73616c IP=0x007d3bc0 Task='mSess1' Task ID=0xc8c2920 fp:  
0x6b61736a sp:0x0c8c1b30 lr:0x007d3b28.
```

- **MSTP Enhancement (PR_1000310463)** - Implemented new CLI command “spanning-tree legacy-path-cost”. See [“MSTP Default Path Cost Controls”](#) on page 38 for details.

Release M.08.91

Problems Resolved in Release M.08.91 (Never released)

- **MSTP Enhancement (PR_1000313986)** - Implemented new CLI command, "spanning-tree legacy-mode".
- **RADIUS (PR_1000316158)** - After a switch reboot, the switch does not recognize a response from a RADIUS or TACACS server.
- **Performance Enhancement (PR_1000291806)** - Allow user configuration of the packet buffer queuing mode. For details, see [“QoS Pass-Through Mode”](#) on page 39.

Release M.08.92

Problems Resolved in Release M.08.92 (Not a general release)

- **Config (PR_1000298146)** — Enabling QoS pass-through Mode causes incorrect information to be displayed in the "show configuration" command.

Release M.08.93

Problems Resolved in Release M.08.93 (Not a general release)

- **Help (PR_1000317711)** — In the VLAN menu Help text, the word 'default' is spelled incorrectly.
- **RSTP (PR_1000307278)** — Replacing an 802.1D bridge device with an end node (non-STP device) on the same Switch port, can result in the RSTP Switch sending TCNs.
- **SNMP (PR_1000315054)**— SNMP security violations appear in syslog after a valid SNMPv3 “get” operation.

Release M.08.94

Problems Resolved in Release M.08.94 (Not a general release)

- **Enhancements (PR_1000319920)** — Added support for following features:
 - DHCP Option 82 functionality, and
 - UDP broadcast forwarding
- **Menu (PR_1000318531)** — When using the Menu interface, the Switch hostname may be displayed incorrectly.

Release M.08.95

Problems Resolved in Release M.08.95 (Not a general release)

- **STP/RSTP/MSTP (PR_1000300623)** — In some cases STP/RSTP/MSTP may allow a loop, resulting in a broadcast storm.

Release M.08.96

Problems Resolved in Release M.08.96 (Never released)

- **Counters (PR_1000321097)** — Drop counters are displaying incorrect information.
- **Enhancement (PR_1000242392)** — Enabled login "Message of the Day" (MOTD) banner. For details on using this feature, refer to “Custom Login Banners for the Console and Web Browser Interfaces” in Chapter 2 of the *Management and Configuration Guide* for 3400cl and 6400cl switches.
- **Web UI Enhancement (PR_1000290489)** — Enhancement to display Port Name along with Port number on the Web User Interface Status and Configuration screens.

Release M.08.97

Problems Resolved in Release M.08.97 (Never released)

- **OSPF (PR_1000319678)** — Switch does not accept IP fragmented OSPF packets.

Release M.10.01

Note: The M.10.xx software releases run only on the ProCurve 3400cl series.

Problems Resolved in Release M.10.01 (Not a general release)

- **Boot ROM/X-Modem (PR_1000327175)** - Boot ROM I.08.11 allows larger file images to be loaded into flash and corrects Console port (X-Modem) reliability issues.
Note: The first time the 3400cl switch boots up with software version M.10.01 or later, Boot ROM version I.08.11 is automatically installed.
- **Crash/ACL (PR_1000323675)** — The Switch may crash with a message similar to:
`ASSERT: Software exception at aaa8021x_proto.c:501 -- in 'm8021xCtrl'`
- **ICMP (PR_1000235905)** — Switch does not send a 'destination unreachable' response message when trying to access an invalid UDP port.
- **SNMPv3 (PR_1000325021)** — Under some conditions, SNMPv3 lines are not written to the running-configuration file.

Release M.10.02

Problems Resolved in Release M.10.02 (Not a general release)

- **Enhancement (PR_1000328392)** — Added RADIUS assigned ACLs.
- **Enhancement (PR_1000328716)** — Added new "show sFlow" commands.

Release M.10.03

Problems Resolved in Release M.10.03 (Never released)

- **Crash/sFlow(PR_1000322009)** — The Switch may crash with a message similar to:
`Software exception in ISR at queues.c:123.`
- **Crash/sFlow (PR_1000327132)** — The Switch may crash with a message similar to:
`Software exception in ISR at btmDmaApi.c:304.`

- **sFlow (PR_1000321195)**— A network management application may incorrectly report spikes in traffic when sFlow is first re-enabled.

Release M.10.04

Problems Resolved in Release M.10.04 (Never released)

- **Enhancement (PR_1000330743)** — Denial of Service logging enhancement with implementation of Instrumentation Monitor. See [“Instrumentation Monitor” on page 70](#) for details.
- **Enhancement (PR_1000331027)** — TCP/UDP port closure feature added. See [“TCP/UDP Port Closure” on page 75](#)
- **STP/RSTP/MSTP (PR_1000330532)** — Improved the "show" commands display of STP port detail information to assist in monitoring and troubleshooting of the spanning tree protocol. See [“Spanning Tree Show Commands” on page 77](#) for details.

Release M.10.05

Problems Resolved in Release M.10.05 (Not a general release)

- **Enhancement (PR_1000311510)** — Ping conformance as defined in RFC 2925.
- **SSHv2 (PR_1000320822)** — The Switch does not generate SSHv2 keys and may crash with a message similar to:

```
TLB Miss: Virtual Addr=0x00000000 IP=0x80593a30 Task='swInitTask' Task
ID=0x821ae330 fp:0x00000000 sp:0x821adfb8 ra:0x800803f0 sr:0x1000fc01.
```

Release M.10.06

Problems Resolved in Release M.10.06

- **CLI (PR_1000334412)** — Operator can save manager config changes.
- **Crash/STP (PR_1000335117)** — Improvement of the PR_1000300623 fix, first included in M.08.95.
- **Enhancement (PR_1000330704)** — RADIUS Command Authorization and Accounting for the Command Line Interface.
- **Log (PR_1000323790)** — Non-ProCurve mini-GBICs identified, but logged only as "self test failure" instead of “unsupported”.
- **OSPF (PR_1000323201)** — OSPF with MD5 does not always redistribute connected networks.

Software Fixes in Release M.08.51 - M.10.72

Release M.10.07

- **Stacking (PR_1000311510)** — When stacking is enabled, a stack member cannot be 'pinged' using the stack number.
- **STP (PR_1000335141)** — The output of the 'show span' CLI command displays a numeral in the 'Type' column, as opposed to terms such as "10/100T".
- **Enhancement (PR_1000309540)** — Added support for the J8440B 10-GbE X2-CX4 Transceiver.
- **Web (PR_1000302713)** — When using the web interface and a large amount of stacking interactions occur, portions of the information from the stack commander may no longer appear.

Release M.10.07

Problems Resolved in Release M.10.07

- **Crash (PR_1000335747)** — Execution of 'configtest' test mode command causes switch to crash with a message similar to:

```
Software exception at parser.c:7898. in 'mSess1', task ID = 0x16726c0  
-> ASSERT: failed. Support: This is a test mode command.
```
- **Enhancement (PR_1000340595)** — Added support for PIM Dense Mode. For details, refer to Chapter 5, “PIM-DM (Dense Mode) on the 5300xl Switches” in the *Advanced Traffic Management Guide for the ProCurve Series 6400cl/5300xl/4200vl/3400cl Switches*.
- **Menu (PR_1000319651)** — The Save option on the "Internet (IP) Service" menu screen not working.
- **Ping MIB (PR_1000311510)** — If the DNS hostname given to ping was invalid (for example hp.com) the switch will crash with an “ASSERT in ip_util.c”.
- **Transceiver (PR_1000310852)** — 10gig LR port has excessive link toggles during bootup.

Release M.10.08

Problems Resolved in Release M.10.08

- **CLI (PR_1000330553)** — Garbage characters displayed in "show snmp-server" cli output.

Release M.10.09

Problems Resolved in Release M.10.09

- **CLI (PR_1000317554)** — The show version command does not display full minor version if it's three digits.
- **Counters (PR_1000327308)** — 10gig port in xSTP blocking mode will increment RX drops on broadcast packets.
- **DHCP (PR_1000343149)** — Client cannot obtain an IP address when two DHCP servers are connected on different local networks
- **Enhancement (PR_1000344652)** — Unidirectional Fiber Break Detection enhancement. See [“Uni-Directional Link Detection \(UDLD\)” on page 80](#) for details
- **SNMPv3 Enhancement (PR_1000338847)** — Added support for the Advanced Encryption Standard (AES) privacy protocol for SNMPv3.
- **VLAN (PR_1000284852)** — Switch transmits packets with VLAN ID 4095.

Release M.10.10

Problems Resolved in Release M.10.10

- **Boot ROM (PR_1000341706)** — Downloading software image results in Flash Boot error:

```
"Bad code in FLASH. Flash memory needs reprogramming or chassis could be faulty. Use a PC as the console and perform the update procedure from the backup floppy diskette. If unsuccessful w/ downloading, then try replacing chassis."
```
- **CLI (PR_1000334494)** — In the “show vlans” command, the “VLAN ID” field is blank.
- **Enhancement (PR_1000336169)** — Added support for STP Per Port BPDU Filtering and SNMP Traps. See [“Spanning Tree Per-Port BPDU Filtering” on page 88](#) for details.
- **Ping MIB (PR_1000337818)** — The handling of multiple ping probe requests is changed such that the requests are sent out one by one instead of being sent all at once. If the pingCtlRowStatus is set to NotInService, the entries of the pingResultsTable and the ping-ProbeHistoryTable get freed. The pingCtlRowStatus cannot be set to NotInService when the pingResultsOperStatus is enabled.
- **Web-UI (PR_1000340311)** — When using the web user interface and accessing the “Security” tab, the switch will request the manager username and password. Then select the “Port Access” button, a second log-in box appears and requests the same manager username and password multiple times, causing the IE browser to hang and requiring the browser to be reset.

Release M.10.11

Problems Resolved in Release M.10.11

- **Crash (PR_1000336436)** — A “get/put” operation on config file via SCP crashes the box with an error message similar to:

```
Software exception at ssh_alarm.c:304 -- in 'mSshAlrm', task ID =  
0x6132588 -> ASSERT: failed.
```

- **Transceiver (PR_1000349320)** — CX4 ports lose configs; "show int config" shows an empty slot rather than CX4.

Release M.10.12

Problems Resolved in Release M.10.12

- **Crash (PR_1000351261)** — On bootup, the switch with a fixed CX4 card installed will crash with the following message:

```
Software exception at gamma_xcvr_util.c:1018. Support: this fix is  
QA only.
```

- **Crash (PR_1000348454)** — Crash when a loop is formed on the network, with error message:

```
NMI event SW:IP=0x002030b4 MSR:0x0000b032 LR:0x002030d4 Task='mMst-  
pCtrl' Task ID=0x60d6060cr: 0x48000040 sp:0x060d5cc8xer:0x00000000
```

- **Crash (PR_1000350363)** — Switch crashes when pinging any other HP switch that is being rebooted, with the following message:

```
Software exception at cli_oper_action.c:986 -- in 'mSess1', task  
ID = 0x62ff180 -> ASSERT: failed
```

- **Radius EAP (PR_1000334731)** — PEAP/TLS Eap Types with IAS Radius Server fail to authenticate.

Release M.10.13

Problems Resolved in Release M.10.13

- **Crash (PR_1000352922)** — The switch may crash with a message similar to

```
Software exception at mstp_ptx_sm.c:118 -- in 'mMstpCtrl', task  
ID = 0x8899e70.-> ASSERT: failed
```

- **Enhancement (PR_1000354065)** — DHCP Protection enhancement for switch 3400.

Release M.10.14

Problems Resolved in Release M.10.14

- **CLI (PR_1000342461)** — Command “show lldp info remote <port number>” reports incorrect information for remote management address.
- **LACP (PR_1000352012)** — LACP state change does not properly reset 10Gig port. Communication through port fails until the port is toggled.
- **LLDP (PR_1000310666)** — The 'show lldp' command does not display information learned from CDPv2 packets.
- **Trunking (PR_1000352851)** — Source Port Filtering on trunks does not work, even though the switch accepts the configuration.
- **XRRP (PR_1000350110)** — XRRP loses layer 3 functionality (pinging) after VLAN is added.

Release M.10.15

Problems Resolved in Release M.10.15

- **CLI (PR_1000358129)** — CLI hangs after running RMON traps code.
- **Crash (PR_1000351410)** — Bus error when pinging switch IP from local serial console.

```
PPC Bus Error exception vector 0x300: Stack-frame=0x067d40e8 HW  
Addr=0x33cc33d2 IP=0x0056a8f8 Task='tNetTask' Task ID=0x67d4278 fp:  
0x00000014 sp:0x067d41a8 lr:
```
- **Crash (PR_1000352177)** — Switch crash when pinging an unreachable host repeatedly, with a message similar to:

```
Software exception at alloc_free.c:362 -- in 'mLinkTest',  
task ID = 0x5be24d0.
```
- **Hang (PR_1000346328)** — Switch hangs during initialization, switch may fail to boot. RMON alarms/events configuration files may be corrupted.

Release M.10.16

Problems Resolved in Release M.10.16 (never released)

- **802.1x (PR_1000353479)** — Changing the supplicant start period (e.g., "aaa port-access supplicant A1 start-period 15") corrupts the supplicant password on a switch that is configured as a supplicant.

Software Fixes in Release M.08.51 - M.10.72

Release M.10.17

- **DHCP Protection (PR_1000360273)** — DHCP Lease renewal packets received on an untrusted port are dropped.
- **DHCP Protection (PR_1000360254)** — An entry with an expired lease is not removed from the binding table.
- **Link Failure (PR_1000361488)** — The J8440B version 10-GbE X2-CX4 may not initialize correctly, causing link failure.
- **Selftest Failure (PR_1000360970)** — A 10-GbE CX4 module (J8434A) will fail selftest following power cycle or software update if it is connected to another switch that is running spanning-tree.

Release M.10.17

Problems Resolved in Release M.10.17

- **Crash (PR_1000367036)** — When a transceiver or mini-GBIC is hot-swapped the switch may crash with a message similar to the following.

```
Software exception at buffers.c:2238 -- in 'mPpmgrCtrl',  
task ID = 0x6351358 -> ASSERT: failed
```

- **Enhancement (PR_1000346164)** — RSTP/MSTP BPDU Protection enhancement. See [“Spanning Tree BPDU Protection” on page 91](#) for details.

Release M.10.18 - Release M.10.19

Releases M.10.18 and M.10.19 were never built.

Release M.10.20

Problems Resolved in Release M.10.20

- **10-GbE <no PR>** — Resolution for failure to initialize the 10-GbE link in port 26 of the ProCurve 3400cl-24G switch, or port 50 of the ProCurve 3400cl-48G switch, after update to software version M.10.17. See [“Known Issues” on page 24](#) for additional information.
- **Enhancement (PR_1000355089)** — This enhancement increases the maximum number of 802.1X users per port to 8.
- **Enhancement (PR_1000355877)** — 802.1X Controlled Directions enhancement. With this change, Administrators can use “Wake-on-LAN” with computers that are connected to ports configured for 802.1X authentication.

- **Enhancement (PR_1000358900)** — A RADIUS accounting enhancement was made. More information about this enhancement will be made available in a future update.

Release M.10.21

Problems Resolved in Release M.10.21 (Not a general release)

- **Crash (PR_1000368540)** — The switch may crash with a message similar to:

```
Software exception at parser.c:8012 -- in 'mSess2',  
task ID = 0x90e10e0 -> ASSERT: failed.
```
- **Crash (PR_1000372183)** — When a meshed network of switches is connected to a non-meshed switch, the meshed switch may crash with a message similar to:

```
Software exception at ldbal_util.c:5970 -- in 'eDrvPoll',  
task ID = 0x5f765a0 -> ASSERT: failed
```
- **Crash (PR_1000376546)** — Rebooting with IGMP enabled may cause the switch to crash with a message similar to:

```
Software exception at sw_sem.c:112 -- in 'swInitTask', task ID =  
0x836b2c40 -> semTake() NULL semaphore: ip_igmp_init.c:1304.
```
- **Web-UI (PR_1000373711)** — Attempting to access the WebUI of a stack member without being logged on as Manager returns a "404 Page Not Found" error.
- **XRRP (PR_1000368594)** — When XRRP infinite failback is enabled, the switch fails to forward packets after a reboot of the Master.

Release M.10.22

Problems Resolved in Release M.10.22 (Not a general release)

- **Crash (PR_1000375501)** — When a link is disconnected and reconnected on a tagged 802.1X supplicant port, the switch may crash with a message similar to:

```
Software exception at macaddr.c:215 -- in 'm8021xCtrl1',  
task ID = 0x8ad9960 -> ASSERT: failed
```
- **Enhancement (PR_1000376406)** — Loop Protection feature additions, including packet authentication, loop detected trap, and receiver port configuration.

Release M.10.23

Problems Resolved in Release M.10.23 (Never released)

- **Crash (PR_1000362248)** — While attempting to configure "qos type-of-service diff-services" the switch may crash with a message similar to:

```
Assertion failed: !VALUE_TOO_BIG_FOR_FIELD, file drvmem.c, line 184.
```

- **Crash (PR_1000378815)** — With a large configuration (>100 VLANs and multiple trunks) and heavy multicast/broadcast traffic, the switch may crash with a message similar to:

```
NMI event SW:IP=0x00643150 MSR:0x0000b032 LR:0x00642320  
Task='mAdMgrCtrl' Task 0 cr: 0x20400020 sp:0x065356e0  
xer:0x20000000.
```

- **Enhancement (PR_1000379804)** — Historical information about MAC addresses that have been moved has been added to the "show tech" command output.
- **Syslog (PR_1000379802)** — Forwarding of event log message to a configured syslog server is not disabled when a specific event log message has been disabled via MIB.

Release M.10.24

Problems Resolved in Release M.10.24 (Never released)

- **CLI (PR_1000364628)** — The command output from "show ip rip peer" yields an improperly formatted peer IP address.
- **Enhancement (PR_1000335860)** — This enhancement provides a configuration option for the source IP address field of SNMP response and generated trap PDUs.
- **Web/RADIUS (PR_1000368520)** — Web Authentication doesn't authenticate clients due to a failure to send RADIUS requests to the configured server.

Release M.10.25

Problems Resolved in Release M.10.25 (Never released)

- **Console/Telnet Hang (PR_1000384178)** — Switch management becomes unresponsive as a result of executing "show int" repeatedly.
- **Enhancement (PR_1000385565)** — (CLI) The port security MAC address limit per port has been increased from 8 to 32 when learn mode is 'static' or 'configured'. However, the global limit of static/configured MAC addresses per ProCurve Series 3400 switch is 400.

- **STP/RSTP/MSTP (PR_1000386113)** — In some cases STP/RSTP/MSTP may allow a loop on 10-Gig ports, resulting in a broadcast storm.

Release M.10.26

Problems Resolved in Release M.10.26 (Not a general release)

- **Enhancement (PR_1000381681)** — This enhancement added eavesdrop protection - the ability to filter unknown Destination IP Address (DA) traffic.
- **MSTP (PR_1000385573)** — MSTP instability when root switch priority is changed. This causes other switches with better priority to each assert themselves to be root thus causing a root war to occur.
- **SNMP (PR_1000388175)** — SNMP PDU configuration enhancement CLI commands are not working.

Release M.10.27

Problems Resolved in Release M.10.27 (Never released)

- **Crash (PR_1000382962)** — Executing the CLI command, "sho int" on a mini-GBIC that isn't linked, may cause the switch to crash with a message similar to:

```
Divide by Zero Error: IP=0x8017becc Task='mSess1' Task ID=0x834b19d0  
fp:0x00000018 sp:0x834b0d20 ra:0x8017be18 sr:0x1000fc01 Division by  
0 Crash at cli_opershow_action.c:1298.
```
- **CLI (PR_1000380660)** — The "show tech transceivers" CLI command displays the wrong message when inserting an "A" version transceiver into a switch that only supports "B" version transceivers. Also, "B" version CX4 transceivers show up as "A" and "A" version SR, LR, and ER transceivers show up as "B" versions.
- **CLI (PR_1000390042)** — Corrupted Spanning Tree Status/Configuration Menu screens.
- **Enhancement (PR_1000374085)** — This enhancement expands the use of the Controlled Directions parameter to also support MAC/Web authentication.
- **MSTP/VLAN (PR_1000381648)** — When a client port is reassigned to a VLAN associated with another MSTP instance, the MAC appears to be incorrectly recorded on the wrong port after that port is assigned back to the original VLAN associated with the other MSTP instance.

Release M.10.28

Problems Resolved in Release M.10.28 (Not a general release)

- **CLI/LLDP (PR_1000377191)** — Output from the CLI command, "show lldp info remote-device <port>" shows a blank field for the chassis ID.
- **CLI (PR_1000390970)** — The command "tftp-enable" is removed from the CLI since that functionality is served by "tftp server/client".
- **CLI/Counters (PR_1000379222)** — Jumbo sized frames received on 10GbE ports decrement the "Total Rx Errors" counters. The 32 bit counter rolls from 0 backwards to 4,294,967,295 and continues to decrement with each received Jumbo frame.
- **Trunking (PR_1000238829)** — Trunks numbered trk10 and greater cause the output from the CLI command "show span" output to be misaligned.

Release M.10.29

Problems Resolved in Release M.10.29 (Never released)

- **CLI/Config (PR_1000375830)** — When using the "no VLAN" command, the user is asked if they want to remove the VLAN. Answering "no" will result in the VLAN being removed anyway.
- **CLI/config (PR_1000391119)** — Copying a configuration file to a switch with a BPDU protection timeout value set may produce an error similar to:

```
CCCCCline: 10007. 1200: Error setting configuration
```
- **CLI (PR_1000390385)** — The CLI help text for "span bpd-protect-timeout" is incorrect; it erroneously displays the help text for "span hello-time".
- **Enhancement (PR_1000376626)** — Enhance CLI "qos dscp-map he" help and "show dscp-map" text to warn the user that inbound classification based on DSCP codepoints only occurs if "qos type-of-service diff-services" is also configured.
- **Lockup (PR_1000394749)** — Switch may lockup when a certain J4858A transceiver type is inserted.
- **SNMP (PR_1000392847)** — RMON alarms that monitor port-specific OIDs are lost if the switch is rebooted.
- **Traceroute (PR_1000379199)** — The reported "traceroute" time is inaccurate; it is one decimal place off.

- **Transceiver hotswap (PR_1000390888)** — Transceiver hotswap issues:
 - Simultaneous hotswap of transceivers on both dual-personality ports will only detect a single change.
 - After certain transceiver hotswaps, the in/out LED indicator will not match the current status of the transceiver.
 - Unsupported mIni-GBIC's hotswapped out of dual personality ports will leave the transceiver in an unknown state of partially inserted.
- **Transceiver hotswap (PR_1000294081)** — The hotswap of a J4858A or B revision wire release style mini-GBIC will result in the switch indicating a port fault condition for that port.
- **Web UI (PR_1000326265)** — Attempting to access the Web UI of a stack member hangs the browser.

Release M.10.30

Problems Resolved in Release M.10.30

- **Daylight savings (PR_1000364740)** — Due to the passage of the Energy Policy Act of 2005, Pub. L. no. 109-58, 119 Stat 594 (2005), starting in March 2007 daylight time in the United States will begin on the second Sunday in March and end on the first Sunday in November.
- **CLI (PR_1000395256)** — The "loop-protect PORT-LIST receiver-action <action>" command does not enable the ports as it should.

Release M.10.31

Problems Resolved in Release M.10.31

- **CLI (PR_1000240838)** — If an invalid time is entered using "clock set" command, the switch responds with an "invalid date" error.
- **CLI (PR_1000199785)** — The tab help function (command-completion) for "IP RIP authentication" is inaccurate. The help selection lists "OCTET-STR Set authentication key" when it should be "ASCII-STR Set RIP authentication key (maximum 16 characters)."
- **Crash (PR_1000398315)** — Under certain conditions when Web Auth is in use, the switch may crash with a message similar to:

```
PPC Bus Error exception vector 0x300:  
  
Stack-frame=0x017b0dd0 HW Addr=0x8200a225 IP=0x00508ce4  
Task='tHttppd' Task ID=0x17b0fa8
```

- **RIP (PR_1000393366)** — The switch does not process RIP (v2) responses containing subnets with a classful subnet mask, when the receiving RIP switch has a connected VLSM network defined that would fall within that classful range.
- **Enhancement (PR_1000372989)** — This enhancement enables the user to set the operator/manager username/password via SNMP.

Release M.10.32

Problems Resolved in Release M.10.32

- **CLI (PR_1000373443)** — The CLI "update" command help text and confirmation message is misleading and confusing.
- **Enhancement (PR_1000376626)** — Enhanced the CLI "qos dscp-map he" help and "show dscp-map" text to warn user that inbound classification based on DSCP codepoints only occurs if "qos type-of-service diff-services" is also configured.
- **Security (PR_1000401384)** — The intrusion flag never comes up for secure ports.
- **RX counters (PR_1000401065)** — ACL *deny* matches on a port cause the Rx Drop counter to increment on software versions M.10.20 or higher.
- **RX counters (PR_1000401395)** — Drops Rx (ifInDiscards) incorrectly increments if a port is blocked by LACP, or if the port receives tagged traffic from a VID for which that port is not a member.
- **Crash (PR_1000392863)** — The switch may crash when "setmib tcpConnState" is used, with a message similar to:

```
NMI event SW:IP=0x0079f4a0 MSR:0x00029210 LR:0x006dca60
Task='eTelnetd' Task ID=0x8a7cbb0 cr: 0x20000042 sp:0x08a7c872
```
- **802.1p (PR_1000392900)** — The switch adds 802.1p Priority 4 to frames forwarded on VLAN tagged ports destined to the IP multicast group 224.0.0.1 (all hosts).
- **RSTP (PR_1000401394)** — When a dynamic LACP trunk transitions to either link-up, or link-down, this action occasionally triggers RSTP instability within the switch. This can result in loops and broadcast storms.
- **Enhancement (PR_1000401306)** — Reload "IN/AT" special enhancement.

Release M.10.33

Problems Resolved in Release M.10.33

- **Crash (PR_1000407542)** — Attempting to change the spanning-tree protocol version from STP to RSTP or MSTP may cause the switch to crash with a message similar to:

```
PPC Bus Error exception vector 0x300: Stack-frame=0x063d5de0 HW  
Addr=0x4b5a697c IP=0x0064c648 Task='mSnmpCtrl'
```

- **QoS (PR_1000370895)** — Once the maximum number of QoS resources is reached, it cannot be cleared without a reboot. The CLI warning message, “Unable to add this QoS rule. Maximum number already reached.” will continue to be displayed until the switch is rebooted.
- **Enhancement (PR_1000408960)** — RADIUS-Assigned GVRP VLANs.
- **DHCP Snooping (PR_1000392148)** — Repeatedly toggling DHCP Snooping on and off may crash the switch with a message similar to: Software exception at bcmHwDsnoop.c:

```
195 -- in 'mAdMgrCtrl', task ID = 0x65a3370 -> BCM ASIC call failed:  
Table full.
```
- **DHCP Snooping (PR_1000403133)** — DHCP-Snooping stops working after some period of time.

Release M.10.34

Problems Resolved in Release M.10.34

- **QoS (PR_1000399873)** — The QoS priority bits are incorrectly set to priority zero on fragmented frames.
- **Menu (PR_1000392862)** — The menu will allow invalid values (greater than 720 sec) to be entered for the Sntp poll interval.
- **Crash (PR_1000410959)** — If the snmpv3 user is deleted on the switch without deleting the associated parameters, then the switch is rebooted, it will repeatedly crash with a message similar to:

```
Software exception at exception.c:373 -- in 'mSnmpEvt', task ID =  
0x17d1818 -> Memory system error at 0x17c22e0 - memPartFree
```
- **Enhancement (PR_1000412747)** — TACACS+ Single Sign-on for Administrators

Release M.10.35

Problems Resolved in Release M.10.35

- **RSTP (PR_1000405368)** — When a primary link goes down and then comes back online, traffic continues on the redundant link and does not shift back to the primary link.

- **BPDU Protection (PR_1000395569)** — BPDU-protection fails after module hot-swap.
- **Enhancement (PR_1000419928)** — The Dynamic ARP Protection feature was added.
- **IP Connectivity (PR_1000418378)** — The switch incorrectly updates its ARP table when a client that is configured with a valid IP address for a valid VLAN, is connected to a port in another VLAN on the switch. This will result in loss of connectivity for the valid client in the appropriate VLAN.

Release M.10.36

Problems Resolved in Release M.10.36 (Never released)

Releases M.10.36 was never built.

Release M.10.37

Problems Resolved in Release M.10.37

- **sFlow (PR_1000396889)** — If the sflow skip count is set greater than the maximum skip count or less than minimum skip count, the switch returns an error, preventing PCM from collecting sampling data.
- **Enhancement (PR_1000369492)** — Update of the MSTP implementation to the latest IEEE P802.1Q-REV/D5.0 specifications to stay in sync with the protocol evolution.

Release M.10.38

Problems Resolved in Release M.10.38 (Not a General Release)

- **TFTP (PR_1000426821)** — TFTP transfers do not work when there is no IP address configured for VLAN 1.
- **PIM-DM (PR_1000398231)** — PIM Dense Mode can trigger unicast routing issues on local hosts.
- **Enhancement (PR_1000428642)** — SNMP v2c describes two different notification-type PDUs: traps and informs. Prior to this software release, only the traps sub-type was supported. This enhancement adds support for informs.
- **SNMP (PR_1000406398)** — URL embedded SNMP traps are not sent as SSL (https) when SSL is enabled, but are sent as plain-text (http) instead. This may result in the trap receiver (e.g. PCM) not being able to display the URL if SSL is enabled.

Release M.10.39

Problems Resolved in Release M.10.39

- **Enhancement (PR_1000428213)** — This software enhancement adds the ability to configure a secondary authentication method to be used when the RADIUS server is unavailable for the primary port-access method.
- **Enhancement (PR_1000415155)** — The ARP age timer was enhanced from the previous limit of 240 minutes to allow for configuration of values up to 1440 minutes (24 hours) or "infinite" (99,999,999 seconds or 3.2 years).
- **Web UI (PR_1000414459)** — During configuration of the GVRP Mode via the web interface (Configuration -> VLAN Configuration -> GVRP Mode), the port list does not show the last three port entries.

Release M.10.40

Problems Resolved in Release M.10.40

- **Web UI (PR_1000380278)** — The switch may periodically require reboot in order for the Web UI authentication page to load.
- **CLI (PR_1000438486)** — When using the **port-access mac-based** CLI command, the client MAC address is sent, in lower case, as the username to the RADIUS server. This fix adds an option so that the MAC address is in uppercase when sent to the RADIUS server. This fix adds additional parameters to the CLI command to support this: "aaa port-access mac-based addr-format".

Release M.10.41

Problems Resolved in Release M.10.41

- **ARP Protection (PR_1000438129)** — ARP and ARP protection data may not display correctly following CLI or SNMP status query.
- **802.1x (PR_1000446227)** — Switch 802.1X authentication running over PAP does not work if RADIUS *message authenticator attribute* is required.
- **LACP/MSTP (PR_1000436184)** — When a trunk comes up, only those physical ports that are not in the trunk negotiation state at the time become unblocked. Therefore, using multiple LACP trunks with MSTP may cause loss of network connectivity.
- **AAA/CLI (PR_1000445886)** — This changes the syntax of 'aaa authentication <port-access | mac-based | web-based>' commands which were previously added in PR_1000438486.

- **SCP (PR_1000428142)** — The switch does not exit a secure copy protocol (SCP) session properly.

Release M.10.42

No Problems Resolved in Release M.10.42 (Never Released)

Release M.10.43

Problems Resolved in Release M.10.43 (Never Released)

- **CLI (PR_1000413734)** — MDI/MDIX information shows "N/A" in the CLI output of the command show int brief. It should show either MDI or MDIX.
- **Enhancement (PR_1000428642)** — SNMP v2c describes two different notification-type PDUs: traps and informs. Prior to this software release, only the traps sub-type was supported. This enhancement adds support for informs. For more information, see [“Release M.10.43 Enhancements” on page 126](#).
- **Enhancement (PR_1000452407)** — The Dynamic IP Lockdown feature was added for the 3400cl series switches. For more information, see [“Release M.10.43 Enhancements” on page 126](#).

Release M.10.44

Problems Resolved in Release M.10.44 (Not a Public Release)

- **Loop Protection (PR_1000447746)** — Client-based AAA stops any packets with unauthenticated source MAC-addresses, including BPDU's and loop-protect packets, creating loops that can be hard to detect.
- **Crash (PR_1000456340)** — The switch may crash with a message similar to the following.

```
No message buffers : alloc_free.c:435
```
- **OSPF (PR_1000453794)** — Removing an IP address from a multinetted VLAN, or removing the entire VLAN, causes the switch to stop seeing it's neighbors "Hello" packets, and, ultimately, lose the OSPF adjacency.
- **Boot Image (PR_1000451000)** — Canceling boot system flash <pri | sec> sets the default boot image. Issuing a reload after canceling causes the switch to boot into the canceled flash image.

Release M.10.45

Problems Resolved in Release M.10.45 (Not a Public Release)

- **Web-UI (PR_1000416955)** — Inserting an LH GBIC into dual personality ports results in the LH ports not appearing in the device view.
- **Meshing (PR_1000453201)** — Concurrent use of meshing and spanning tree may result in instability in spanning tree, with chronic root bridge transitions every 20 to 40 seconds.
- **Config (PR_1000400244)** — The switch prompts the user to save config, even though no apparent changes have been made. However, if SNMP sets have occurred in the background, then the user will still see the save config prompt due to the configuration changes caused by the SNMP sets.

Release M.10.46

Problems Resolved in Release M.10.46 (Not a Public Release)

- **SSH (PR_1000453226)** — Configuration of SSH login to the manager mode (using the command, `aaa authentication ssh enable public-key <enter>`) triggers an error “Not legal combination of authentication methods,” but it should be a valid command syntax.
- **DIPLD (1000457808)** — When a user with a DHCP assigned IP address de-authenticates and then re-authenticates, the DIPLD bindings show the port is bound to multiple IP addresses, and the switch will accept traffic from both IP addresses.

Release M.10.47

Problems Resolved in Release M.10.47 (Never Released)

- **SNMP (PR_1000448463)** — The SNMP Engine ID Discovery process described in RFC 3414 is not working properly.
- **SSH/SCP (PR_1000453751)** — The switch does not properly exit a secure copy protocol (SCP) session, particularly when a software image is transferred.
- **Trunking (1000461440)** — When dynamic ARP protection and DHCP snooping are configured, a trunk's trust status cannot be configured from the appropriate interface configuration context.
- **SSH/SCP (PR_1000742969)** — The following issues with using SSH/SCP were fixed.
 - In **show ip ssh**, sessions 3 and 4 may display "console" instead of "inactive," when those sessions are not in use.

- The switch does not send an appropriate exit status message to the client. This corrects the symptom that occurs in some applications, which reports a message similar to:

```
Fatal error: Server unexpectedly closed connection.
```
 - The SSH client application does not get a command prompt (or equivalent) back from the switch until the OS is verified and burned to flash.
 - The **show flash** command incorrectly shows an OS image present in flash before the OS is completely copied to flash.
- **Mirroring (PR_1000460844)** — When multiple VLANs are configured on the same port and VLAN monitoring is enabled, packets to other VLANs are mirrored.

Release M.10.48

Problems Resolved in Release M.10.48 (Not a Public Release)

- **Daylight Savings (PR_1000467724)** — DST is outdated for the Western European time zone. This change corrects the schedule for Western Europe time zone: DST to start the last Sunday in March and DST to end the last Sunday in October.
- **SNMP (PR_1000715545)** — Buffered log messages (those log messages occurring in the switch's event log prior to an IP address being enabled) are not filtered properly at boot up when the switch is configured to send those log messages as traps. For example, non-critical log entries may get sent to trap destinations configured to receive only critical events.
- **Crash (PR_1000464345)** — The characters "IP(" when present as part of a port name may cause config corruption when the switch is restarted.
- **SFTP/SCP (PR_1000428974)** — SFTP or SCP transfer of the configuration files fails to complete.

Release M.10.49

Problems Resolved in Release M.10.49 (Not a Public Release)

- **Port Security (PR_1000755715)** — Port Security Mac Lockdown send-disable allows a few frames through.
- **Running/Startup Config (1000750637)** — SNMPv3 users are not correctly reflected in startup config. The running-config and startup-config are not in sync, after a fresh configuration load and reboot.
- **Sflow (PR_1000749192)** — Trunked ports show the following sFlow sampling inaccuracies:

Routed traffic is off by a factor of 1000
Switched traffic is not sampled at all

- **Security (PR_1000388616)** — Possible cross-site scripting vulnerability in Web Management Interface.
- **Config (PR_1000763386)** — An SNMPv3 user is not reflected in startup config as it should be. This is an additional fix for PR_1000750637.
- **CLI (PR_1000713515)** — When the CLI command **erase startup config** is issued, the switch asks whether you want to save the config.
- **Transceivers (PR_1000467314)** — Revision C mini-GBICs are being incorrectly displayed as revision B.
- **Crash (PR_1000471594)** — Use of the CLI command **show config <file>** may cause the switch to crash with a message similar to the following.

```
TLB Miss: Virtual Addr=0x00000004 IP=0x80150828 Task='mSess1' Task
ID=0x85e48550 fp:0x85e47978 sp:0x85e478e0 ra:0x801507cc
sr:0x1000fc01
```

- **Crash (PR_1000464612)** — Booting from the secondary image, or some types of configuration file manipulation (for example, use of the CLI command **erase start**) may cause the switch to crash with a message similar to the following.
- ```
Software exception at ConfigTree.cc:508 -- in 'mChassCtrl'.
```
- **Tagged/Untagged VLANs (PR\_1000759034)** — Trunks may transmit tagged frames from untagged VLANs.
  - **QoS (PR\_1000454194)** — When using the Web Management interface to configure VLAN based QoS, the VLANs do not show up in the configuration window as selectable when applying a DSCP or QoS value.
  - **Web Auth (PR\_1000380278)** — After running for a period, the switch will get into a state in which it must be rebooted in order for the Web authentication page to load.

## Release M.10.50 through M.10.64

### Problems Resolved in Release M.10.50 - M.10.64 (Never Built)

## Release M.10.65

### Problems Resolved in Release M.10.65 (Not a Public Release)

- **Authentication (PR\_1000454714)** — Concurrent 802.1X and MAC Authentication does not give the 802.1X value precedence. This fix gives 802.1X VLAN assignment precedence over MAC Auth RADIUS VLAN assignment.
- **Web Management (PR\_1000760153)** — A Java error occurs when viewing "Stack Closeup" in the Web Management interface. Only a blank screen is displayed.
- **DHCP Snooping (PR\_1000469934)** — When DHCP Snooping is enabled and configured, and a client sends a "DHCPINFORM" after receiving address information, the DHCP server response is not forwarded to the client by the switch.
- **IDM/DIPLD (PR\_1000784427)** — The IDM feature that allows RADIUS-assigned ACLs for 802.1X clients was incorrectly enabled to support DIPLD CLI commands.
- **RADIUS/Jumbo (PR\_1000779048)** — When an 802.1X enabled port belongs to a VLAN that is jumbo enabled, the Access-Request will specify a value of Framed-MTU of 9182 bytes. When the RADIUS server replies with a large frame, the switch does not respond, causing the authentication process to halt.
- **Crash (PR\_0000001756)** — Some SNMP set commands may cause the switch to crash with a message similar to the following.

```
Software exception at bcmHwVlans.c:149 -- in 'mAdMgrCtrl', task
ID=0x18636e8 -> ASIC call failed: Entry not found.
```
- **Crash (PR\_0000002433)** — A certain sequence of CLI commands may cause the switch to crash with a message similar to the following.

```
Software exception at dsnoop_ctrl.c:109 -- in 'mDsnoop002'.
```
- **Enhancement (PR\_0000001316)** — The MSTP VLAN Assignment is enhanced. For more information, see ["Release M.10.65 Enhancements" on page 136](#).
- **Config (PR\_0000002077)** — Upload of configuration files fails when the spanning-tree CLI command is present in the config file.
- **VLAN (PR\_0000002103)** — The alteration of the VLAN/MSTP instance mapping in the pending configuration is not functioning properly. Any attempt to remove a single VLAN ID (VID) from one MSTP instance and then assign it to another MSTP instance fails, though specifying a VID range succeeds.

## Release M.10.66

### Problems Resolved in Release M.10.66 (Not a Public Release)

- **UDLD (PR\_0000002473)** — UDLD protocol packets received on a (non-UDLD) trunk port are incorrectly forwarded out of the same port from which they are received, resulting in high CPU usage on the switch.

- **CLI (1000415243)** — Output from the CLI command **show name** still lists 10-GbE transceiver names, even after the transceivers are removed and replaced with another type of transceiver.
- **CLI (PR\_1000430534)** — Output from the **show port-access** mac-based CLI command may omit connected clients.
- **Enhancement (0000000818)** — Enhancement to allow syslog configuration via SNMP. For more information, see [“Release M.10.66 Enhancements” on page 140](#).

## Release M.10.67

### Problems Resolved in Release M.10.67 (Never Released)

- **RADIUS Accounting (PR\_0000004145)** — An incomplete "Calling-Station-ID" field is sent in the accounting-request to the RADIUS server on the boot system command.
- **RADIUS Accounting (PR\_0000004141)** — The "Acct-Status-Type" attribute is missing in the accounting-request to RADIUS server on the boot system command.
- **RADIUS Accounting (PR\_0000004139)** — Procurve switches do not send the accounting request to RADIUS server on reload command.
- **TELNET hang (PR\_1000457765)** — If **Ctrl+S** is typed and then the Telnet window is closed, the Telnet session may become unresponsive and fail to reset with the kill command issued at the console prompt. This may require the switch to be reloaded to become active again.
- **GVRP/RADIUS (PR\_0000006051)** — RADIUS assigned VLANs are not correctly propagated in GVRP.
- **Web-Management (PR\_0000002153)** — The Web Management interface does not allow configuration of 'static' port-security using a value higher than 8. Selecting 9-32 will change the value to 'limited' in the drop-down box.
- **MAC Auth/802.1X (PR\_0000004095)** — When an 802.1X session is closed, the client logout is not seen by the switch. Since the switch does not end the 802.1X session, the client does not succeed at MAC Authentication. This issue has been seen when using an 802.1X client such as the Odyssey Access Client Manager, version 4.60.
- **802.1X Port Lock-Up (PR\_0000005372)** — If the first frame is sent from an all zero's MAC address to a broadcast destination address, an 802.1X port will freeze, AAA will quit functioning.
- **Web-Management (PR\_1000451437)** — The [?] button in the upper right corner of the Web Management Interface pages brings up an obsolete text message unless there is a “Management Server URL” configured on the switch.

- **Crash (PR\_0000004023)** — Repeated PCM configuration scans using SSH/SCP may cause the switch to crash with a message similar to the following.

```
PPC Data Storage (Bus Error) exception vector 0x300:
Stack Frame=0x07af44c0 HW Addr=0x6520463a IP=0x00965a88
Task='tSsh0' Task ID=0x7af4810
```

## Release M.10.68

### Problems Resolved in Release M.10.68

- **TACACS+ (PR\_0000003839)** — The TACACS server configuration parameter accepts an address from an invalid/reserved IP range: 0.0.0.1 to 0.255.255.255.
- **MAC Table (PR\_0000005185)** — Hardware MAC table entries are not getting deleted appropriately.
- **Enhancement (PR\_0000003127)** — Link Trap and LACP Global enable/disable feature. For more information, see [“Release M.10.68 Enhancements” on page 143](#).

## Release M.10.69

### Problems Resolved in Release M.10.69 (Not a Public Release)

- **Management (PR\_0000005902)** — The switch management may become unresponsive, resulting in loss of TELNET, Web Management, and console access functionality of the switch.
- **PCM (PR\_0000008113)** — Repeated ProCurve Manager Config Scans may trigger subsequent Config Scan failure.
- **SFTP/SCP (PR\_0000008270)** — Beginning with software version M.10.67, SFTP/SCP will not close the "client" session after the file transfer. The client session will need to be manually closed.
- **PC phone/authentication (PR\_0000008777)** — When using an IP phone in tandem with a PC connected to the phone, the phone will sometimes come up using untagged packets until acquiring its tagged VLAN and priority information. In this case the IP phones untagged MAC address will block the PC communicating to the port until the phone's MAC address expires (default 5 minutes).
- **802.1X (PR\_0000008780)** — 802.1X does not receive expiration notifications from port security if 802.1X is running alone, without Web or MAC-based Authentication.

- **PC Phone/Authentication (PR\_0000007209)** — When an IP phone is used in tandem with a PC connected to the phone, if the phone is moved to a tagged VLAN, some phone manufacturers send some traffic to the switch untagged. This may result in traffic disruption including the PC not being allowed to authenticate.
- **CLI (PR\_0000002815/1000406763)** — Output from the **show tech** CLI command was modified to include output from **show access-list resources** and **show access-list radius all** commands.
- **DHCP Snooping (PR\_0000009387)** — The **max-vlan** configuration may result in a different MAC address being used in switch-generated DHCP requests from non-default VLANs versus the default VLAN. When DHCP-Snooping is globally enabled in this situation, the DHCP server offers are dropped to the non-default VLANs.
- **PC Phone/Authentication (PR\_0000010104)** — When using an IP phone in tandem with a PC, sometimes the VLAN assignment after authentication of the PC is delayed.
- **802.1X (PR\_0000010275)** — For a port that is being authenticated via 802.1X, the user fails authentication if the **unauth vid** value is configured.
- **CLI (PR\_0000010942)** — The CLI command output for **show run** does not display **aaa port-access <port#>** when MAC-based authentication with mixed port access mode is configured. Other **show** commands may also be affected.
- **CLI (PR\_0000010378)** — Session time (sec.) remains at zero in response to the CLI command **show port-access authenticator <port> session-counters**; it should increment.
- **Enhancement (PR\_0000010783b)** — The CLI output for the **show tech transceivers** command has been enhanced to be consistent with other platforms.

## Release M.10.70

### Problems Resolved in Release M.10.70 (Not a Public Release)

- **Config (PR\_0000007953)** — The config line **spanning-tree instance <n> vlan <vid>** is truncated in some cases, causing loss of configuration after reload of the config file.
- **DHCP / DHCP Snooping (PR\_0000008118)** — Switches may be intermittently unable to get an IP address via DHCP after enabling DHCP snooping.
- **Authentication (PR\_0000012553)** — The switch sends EAP supplicant packets with the identity field truncated to 24 bytes after a reload.

- **Dynamic ARP Protection (PR\_0000009942)** — When a switch using Dynamic ARP Protection is rebooted, it blocks all ARP traffic on untrusted ports, including traffic considered valid according to the binding database. On trusted ports, traffic flows normally. Workarounds: either disable / re-enable ARP protect, or configure ports to be trusted, and then untrusted again.
- **802.1X Authentication (PR\_0000011718)** — When an 802.1X enabled port belongs to a VLAN that is jumbo enabled, the Access-Request will specify a value of Framed-MTU of 9182 bytes. This allows the RADIUS server to reply with a large fragment which the switch does not process, causing the authentication to fail.
- **CLI (PR\_0000009868)** — Execution of a **show** command in one telnet or console session prevents successful execution of a **show** command in a concurrent management (CLI) session.
- **Crash (PR\_0000003523)** — Disabling a port with Dynamic IP Lockdown (DIPLD) and 802.1X enabled may cause the switch to reboot unexpectedly with a message similar to the following:

```
Software exception at idmCommonAcl.c:1547 -- in 'midmCtrl', task ID
= 0x85cc1eb0
```

- **802.1X (PR\_0000010850)** — If an **unauth-vid** is configured, and the client limit is reached on a switch port, a properly credentialed re-authentication following an improperly credentialed authentication attempt (for example, incorrect password) will leave the 802.1x client in the unauthorized VLAN instead of applying the appropriate authorized VLAN.
- **SNMP Traps (PR\_0000007448/1000469020)** — The switch no longer sends warm/cold start SNMP traps on reload/boot.
- **802.1X (PR\_0000009344)** — The switch sends an EAP-notification out to the client as a result of the Radius server sending a Reply-Message in the Access-Accept packet after EAP-Success. This fix follows the suggestion in RFC 3579, section 2.6.5 and silently discards attributes sent out after the authentication is complete.
- **Crash (PR\_0000009411)** — A switch with 802.1X and RADIUS accounting configured may experience an unexpected reboot with a message similar to the following.

```
Software exception at aaa8021x_proto.c:255 -- in 'm8021xCtrl'
```

- **CLI Help (PR\_0000010484)** — The CLI tab completion for the command parameter **[ethernet] PORT-LIST** should list the **all** option, but it does not.
- **Crash (PR\_0000012124)** — Switches configured for meshing may reboot unexpectedly with a message similar to the following.

```
Software exception at ldbal_util.c:2525 -- in 'mLdBalCtrl'
```

- **MAC Authentication (PR\_0000011949)** — MAC authentication may fail to take place unless the switch port status is toggled.



## Release M.10.71

### Problems Resolved in Release M.10.71 (Not a Public Release)

- **802.1X (PR\_0000014842)** — If an invalid number of characters are used at the CLI for the command `aaa port-access supplicant <port number> secret`, the CLI returns an error message that references the wrong port number for the supplicant being configured.
- **Enhancement (PR\_0000011636)** — This enhancement adds the client's IP address to the RADIUS accounting packets sent to the RADIUS server by the switch. The IP address of the client is included in the RADIUS accounting packet sent by the switch to the RADIUS server. The client obtains the IP address through DHCP, so DHCP snooping must be enabled for the VLAN of which the client is a member.
- **Mirroring (PR\_0000008008)** — When a port that was previously configured to be monitored to a mirror-port has that monitoring configuration removed, traffic continues to be sent to the mirror port until the switch is reloaded.
- **Crash (PR\_0000006336)** — Copying a configuration that contains the lines **interface all lacp and snmp-server enable traps link-change all** from a TFTP server to the switch may cause an unexpected reboot with a message similar to the following.

```
Software exception at cli_xlate.c:3692 -- in 'mftTask', task ID =
0x5ee17f0.
```

- **SNMP (PR\_0000002764)** — The SNMP MIB object that allows authenticator functionality on a port to be enabled or disabled (**hpicfDot1xPaePortAuth**) can be set to an invalid value.
- **802.1X (PR\_0000012568)** — There may be a problem with a login error message.
- **TACACS (PR\_0000008268)** — If an invalid IP address is configured for the TACACS server, the switch will not allow its removal from the configuration at the CLI.
- **Controlled Direction (PR\_0000009818)** — The switch does not properly enable or edit the controlled direction parameter (in the config line **aaa port-access controlled-direction <in/out/both>**) in the configuration.
- **BPDU-Protection (PR\_0000012541)** — The presence of a trunk group in a switch with STP BPDU-protection configured may trigger the switch to block the wrong port when a BPDU is received.

## Release M.10.72

### Problems Resolved in Release M.10.72

- **Authentication (PR\_0000011917)** — The switch does not recognize the "session-timeout" attribute from a RADIUS server following MAC authentication.

- **Config (PR\_0000005002)** — If a friendly port name uses the characters TRUNK=, then after a reload, all the trunking configuration will have been removed from the configuration.
- **GVRP (PR\_0000012224)** — Changing the GVRP **unknown-vlan** state from 'block' to 'learn' and vice versa stops all GVRP advertisements from that interface until the interface is disabled and then re-enabled.
- **RADIUS Accounting (PR\_0000012487/0000037453)** — The switch doesn't send an accounting-stop when a switch reload closes the session.
- **Meshing/STP (PR\_0000004611)** — Concurrent use of meshing and spanning tree on 10-GbE-CX4 ports may result in instability in spanning tree, with chronic root bridge transitions every 20 to 40 seconds.
- **SNMP (PR\_0000017534)** — SNMP communication may cease after a software update and configuration copy to the switch.
- **DHCP-Snooping/IP Lockdown (PR\_0000013457)** — DHCP-Snooping does not block DHCP offers from an untrusted port as it should when the feature is configured in combination with IP source lockdown.
- **Authentication (PR\_0000014177)** — The switch consumes too many packets during the AAA Authentication and dynamic VLAN assignment, potentially causing clients that require bootP to fail to retrieve their configurations and initialize.
- **Authentication (PR\_0000011138)** — If the RADIUS server becomes unavailable, the **eap-radius authorized** option allows the switch to authenticate devices. If the response time of the RADIUS subsystem is greater than the server-timeout value on the switch or the device supplicant then switch will not be able to authenticate devices, and no warning of this failure will be displayed. This fix triggers the display of the following CLI message.

```
The RADIUS connection timeout must be less than the authentication
server timeout for the switch to authenticate automatically when the
RADIUS server is unavailable.
```

- **DHCP-Snooping (PR\_0000012237)** — The DHCP-Snooping binding table has incorrect lease times when the binding table has been read at boot.
- **DHCP-Snooping (PR\_0000018615)** — DHCP-Snooping may not block appropriately on untrusted ports.
- **RADIUS Accounting (PR\_0000017732)** — RADIUS accounting is incrementing the wrong counter in response to a dropped (invalid) packet from the RADIUS server.
- **DHCP-Snooping (PR\_0000018613)** — When a DHCP client is unable to access the network due to receipt of a DHCP-NACK (expected and desired behavior in that circumstance), there is no indication of why the 'failure' is occurring to help the network administrator understand what is happening. This fix adds log messages similar to the following.

Drop offer from <DHCP server IP address> of <DHCP address offer> because the address is assigned to some other client

Drop request from <MAC address of client requesting an IP address that is already in use> for <IP address requested by client> because the address is assigned to some other client

- **DHCP—Snooping (PR\_0000019155)** — DHCP-Snooping does not correctly identify that a packet is a fragment, and drops UDP Fragments if a hex value of 44 (68 Decimal) is present in the payload where the header is usually located (in a non-fragment).
- **Unauthenticated VLAN (PR\_0000010533)** — The switch allows an inherent configuration conflict; an unauthenticated VLAN (**unauth-vid**) can be configured concurrently for both 802.1X and Web/MAC authentication. This fix will not allow concurrent configuration of an **unauth-vid** for the **aaa port-access authenticator** and **aaa port-access web-based** or **aaa port-access mac-based** functions. Software versions that contain this fix will not allow the this configuration conflict at the CLI. *Existing configurations will be altered by this fix*, and an error will be reported at the switch CLI and event log.

*Best Practice Tip:* 802.1X should not have an unauthenticated VLAN setting when it works concurrently with Web-based or MAC-based authentication if the unauth-period in 802.1X is zero (the default value). Recall that the unauth-period is the time that 802.1X will wait for authentication completion before the client will be authorized on an unauthenticated VLAN. If 802.1X is associated with an unauthenticated VLAN when the unauth-period is zero, Web- or MAC-auth may not get the opportunity to initiate authentication at all if the first packet from the client is an 802.1X packet. Alternatively, if the first packet sent was not 802.1X, Web- or MAC-auth could be initiated before 802.1X places the user in the unauthenticated VLAN and when Web- or MAC-auth completes successfully, it will be awaiting traffic (to enable VLAN assignment) from the client but the traffic will be restricted to the unauthenticated VLAN, and thus the client will remain there.

If a MAC- or Web-based configuration on a port is associated with an unauth-VID, and an attempt is made to configure an unauth-VID for 802.1X (port-access authenticator), the switch with this fix will reject the configuration change with a message similar to one of the following.

Message 1 (when an unauth-vid config is attempted on a port with an existing Web- or MAC-auth unauth-vid):

```
Configuration change denied for port <number>.Only Web or MAC-
authenticator can have unauthenticated VLAN enabled if 802.1X
authenticator is enabled on the same port. Please disable Web and
MAC authentication on this port using the following commands:
```

```
"no aaa port-access web-based <PORT-LIST>" or
```

```
"no aaa port-access mac-based <PORT-LIST>"
```

```
Then you can enable 802.1X authentication with unauthenticated VLAN.
You can re-enable Web and/or MAC authentication after you remove the
unauthenticated VLAN from 802.1X.Note that you can set unauthenti-
cated VLAN for Web or MAC authentication instead.
```

Message 2 (when an unauth-vid config is attempted on a port with an existing 802.1X unauth-vid):

Configuration change denied for port <number>. Only Web or MAC-authenticator can have unauthenticated VLAN enabled if 802.1X authenticator is enabled on the same port. Please remove the unauthenticated VLAN from 802.1X authentication on this port using the following command:

```
"no aaa port-access authenticator <PORT-LIST> unauth-vid"
```

Note that you can set unauthenticated VLAN for Web or MAC authentication instead.

Message 3:

Configuration change denied for port <number>. Only Web or MAC-authenticator can have unauthenticated VLAN enabled if 802.1X authenticator is enabled on the same port. Please use unauthenticated VLAN for Web or MAC authentication instead.

Event log message when the configuration is changed:

```
mgr: Disabled unauthenticated VLAN on port <number> for the 802.1X. Unauthenticated VLAN cannot be simultaneously enabled on both 802.1X and Web or MAC authentication.
```

- **Crash (PR\_0000038448)** — Switches configured for Web Authentication may reboot unexpectedly in response to DHCP activity, displaying a message similar to the following.

```
Software exception at exception.c:621 -- in 'mAcctCtrl', task ID = 0x842d140 -> Memory system error at 0x7ed5950 - memPartFree
```

- **DHCP-Snooping (PR\_0000038432)** — DHCP-Snooping logs an incorrect server ID for ACK packets, impairing the effectiveness of the fix described in [PR\\_0000018613](#).



© 2004 - 2009 Hewlett-Packard Development Company, LP. The information contained herein is subject to change without notice.

October 2009  
Manual Part Number  
5991-4764

## Free Manuals Download Website

<http://myh66.com>

<http://usermanuals.us>

<http://www.somanuals.com>

<http://www.4manuals.cc>

<http://www.manual-lib.com>

<http://www.404manual.com>

<http://www.luxmanual.com>

<http://aubethermostatmanual.com>

Golf course search by state

<http://golfingnear.com>

Email search by domain

<http://emailbydomain.com>

Auto manuals search

<http://auto.somanuals.com>

TV manuals search

<http://tv.somanuals.com>