



Software Update C.09.22 Release Notes

for the HP ProCurve Switches 1600M, 2400M, 2424M, 4000M, and 8000M

Topics:

- Software Management (page [7](#))
- TACACS+ Authentication for Centralized Control of Switch Access Security (page [8](#))
- CDP (page [31](#))
- New Time Synchronization Protocol Options (page [42](#))
- Operation and Enhancements for Multimedia Traffic Control (IGMP) (page [49](#))
- New: Menu Enhancement for Moving from Operator Access to Manager Access (page [60](#))
- Configuring and Using HP ProCurve Stack Management (page [61](#))
- Using the Auto-10 Port Configuration Option (page [85](#))
- Updates to VLAN Configuration Options (page [85](#))
- Enhanced Multimedia Traffic Filtering (page [86](#))
- FAQs from the HP ProCurve web site (page [87](#))
- Release History (page [99](#))

Caution: Archive Pre-C.09.xx Configuration Files

A configuration file saved under version C.09.*xx* software is not backwards-compatible with previous software versions. For this reason, HP recommends that you save a copy of any pre-C.09.*xx* configuration file *before upgrading* to C.09.*xx* or later software releases in case there is ever a need to revert back to pre-C09.*xx* software. Instructions for saving a copy of the current configuration file are found in the "File Transfers" chapter of the *Management and Configuration Guide* provided for the switch.

© Copyright 2001, 2004 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice.

This document contains proprietary information, which is protected by copyright. No part of this document may be photocopied, reproduced, or translated into another language without the prior written consent of Hewlett-Packard.

Publication Number

5969-2375
May 2004
Edition 4

Applicable Products

HP ProCurve Switch 1600M	(J4120A)
HP ProCurve Switch 2400M	(J4122A)
HP ProCurve Switch 2424M	(J4093A)
HP ProCurve Switch 4000M	(J4121A)
HP ProCurve Switch 8000M	(J4110A)

Trademark Credits

Windows NT®, Windows®, and MS Windows® are US registered trademarks of Microsoft Corporation.

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.

Contents

Software Management

Downloading Switch Documentation and Software from the Web	7
--	---

TACACS+ Authentication for Centralized Control of Switch Access Security

Before You Begin	14
Overview	14
Viewing and Configuring Authentication Options	15
Viewing and Configuring the Timeout Period and (Global) Encryption Key	19
Viewing and Configuring TACACS+ Servers and (Per-Server) Encryption Keys	20
General Authentication Process Using a TACACS+ Server	25
Local Authentication Process	26
General Operation	27
Encryption Options in the Switch	27

CDP

Introduction.....	31
Outgoing Packets	32
Incoming CDP Packets	33
Viewing and Changing the Switch's Current CDP Configuration	36

New Time Synchronization Protocol Options

General Steps for Running a Time Protocol on the Switch:	43
Disabling Time Synchronization	43
Viewing and Configuring SNTP	44
Viewing and Configuring TimeP	47

Operation and Enhancements for Multimedia Traffic Control (IGMP)

How Data-Driven IGMP Operates.....	49
New: IGMP Now Operates With or Without IP Addressing.....	50
Configuration Options for Forced Fast-Leave	53
Listing the Forced Fast-Leave Configuration	54

Configuring Per-Port Forced Fast-Leave IGMP	56
Querier Operation	57
Changing the Querier Configuration Setting	58

New: Menu Enhancement for Moving from Operator Access to Manager Access

Configuring and Using HP ProCurve Stack Management

Components of HP ProCurve Stack Management	62
General Operation.	62
Operating Rules for Stacking.	64
Configuring and Bringing Up a Stack.	66
Monitoring Stack Status.	78
SNMP Community Operation in a Stack	82
Stacking Operation with a Tagged VLAN.	83
Status Messages	83
Changes to the Web Browser Interface for Commander Switches	84

Using the Auto-10 Port Configuration Option

Updates to VLAN Configuration Options

Enhanced Multimedia Traffic Filtering

FAQs from the HP ProCurve Web Site

Auto-Negotiation	87
Cabling	88
Features	88
Gigabit Stacking	91
Modules	91
Spanning Tree Protocol	93
Troubleshooting	95
Trunking (HP and Fast EtherChannel—FEC)	95
VLANs	95

Release History

Release C.09.22	99
Release C.09.19	100
Release C.09.18	101
Release C.09.16	101
Release C. 09.14 (BETA)	102
Release C.09.13	103
Release C.09.12 (BETA)	103
Release C.09.10	103
Release C.09.09	104
Release C.09.08 (BETA)	104
Release C.09.07 (BETA)	104
Release C.09.06 (BETA)	104
Release C.09.05 (BETA)	105
Release C.09.04 (BETA)	105
Release C.09.03 (BETA)	106
Release C.09.02	106
Release C.08.23 (BETA)	107
Release C.08.22	108
Release C.08.21 (BETA)	108
Release C.08.20 (BETA)	108
Release C.08.04 (BETA)	108
Release C.08.03	108
Release C.08.02 (BETA)	109
Release C.08.01 (BETA)	109
Release C.07.28 (BETA)	109
Release C.07.27	110
Release C.07.26	110
Release C.07.25	110

Release C.07.24	110
Release C.07.23	111
Release C.07.22	111
Release C.07.21 (BETA)	111
Release C.07.20 (BETA)	112
Release C.07.02	112
Release C.07.01	113
Release C.06.07 (BETA)	114
Release C.06.06	114
Release C.06.05 (BETA)	114
Release C.06.04 (BETA)	115
Release C.06.03	115
Release C.06.02 (BETA)	115
Release C.06.01	116
Release C.05.07	117
Release C.05.06	118
Release C.05.05	118
Release C.05.04	119
Release C.05.03	119
Release C.05.02	119

Software Management

Software Update Notice

Check the HP ProCurve web site frequently for *free* software updates for the various HP ProCurve switches you may have in your network.


Downloading Switch Documentation and Software from the Web

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

To Download a Software Version:

1. Go to HP's ProCurve web site at:
<http://www.hp.com/go/hpprocurve>.
2. Click on **Software updates**.
3. Under **Latest software**, click on **Switches**.

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

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

TACACS+ Authentication for Centralized Control of Switch Access Security

TACACS+ authentication in an HP Switch 1600M, 2400M, 2424M, 4000M, or 8000M enables you to use a central server to allow or deny access to the switch (and other TACACS-aware devices) in your network. This means that you can use a central database to create multiple unique username/password sets with associated privilege levels for use by individuals who have reason to access the switch from either the switch's console port (local access) or Telnet (remote access).

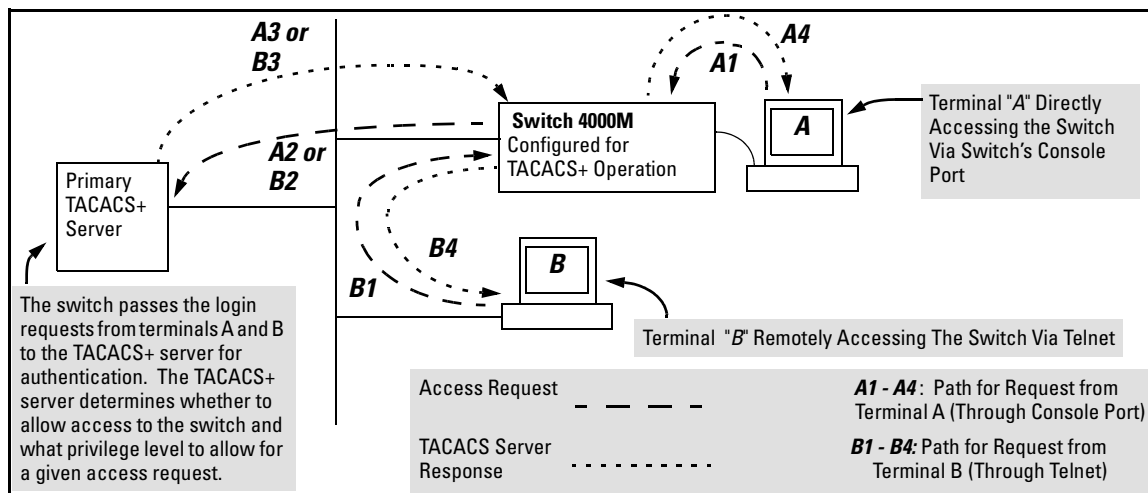


Figure 1. Example of TACACS+ Operation

Overview of Operation. TACACS+ on HP switches uses an authentication hierarchy consisting of remote control through a TACACS+ server and local passwords configured in the switch. That is, with TACACS+ configured on the switch, the switch first tries to contact a designated TACACS+ server for authentication services. If the switch fails to connect to any TACACS+ server, it can default to its own locally assigned passwords for authentication control, if it has been configured to do so.

Notes Regarding Software Release C.09.xx

The HP ProCurve Switches 1600M, 2400M, 2424M, 4000M, and 8000M supports TACACS+ authentication, which allows or denies access to a device on the basis of correct username/password pairs managed by the TACACS+ server. This release does not support TACACS+ authorization or accounting services.

TACACS+ does not affect web browser interface access. To block unauthorized access through the web browser interface, see "Controlling Web Browser Interface Access When Using TACACS+ Authentication" on page 28.

For more on general authentication operation, see “How Authentication Operates” on page 25.

Terminology and Options Used in TACACS Applications:

- **NAS (Network Access Server):** This is an industry term for a TACACS-aware device that communicates with a TACACS server for authentication services. Some other terms you may see in literature describing TACACS operation are *communication server*, *remote access server*, or *terminal server*. These terms apply to an HP switch when TACACS+ is enabled on the switch (that is, when the switch is TACACS-aware).
- **TACACS+ Server:** The server or management station configured as an access control server for TACACS-enabled devices. To use TACACS+ with a TACACS-aware HP switch and any other TACACS-capable devices in your network, you must purchase, install, and configure a TACACS+ server application on a server or management station in the network. The TACACS+ server application you install will provide various options for access control and access notifications. For more on the TACACS+ services available to you, see the documentation provided with the TACACS+ server application you will use.
- **Authentication:** The process for granting user access to a device through entry of a user name and password and comparison of this username/password pair with previously stored username/password data. Authentication also grants levels of access, depending on the privileges assigned to a user name and password pair by a system administrator.
 - **Local Authentication:** This method uses passwords configured locally on the switch; one each for manager-level (read-write) and operator-level (read-only) access to the switch. You can assign local passwords through the Menu or web browser interfaces. (The web browser interface also allows you to assign a local username.) Because this method assigns passwords to the switch instead of to individuals who access the switch, you must distribute the password information on each switch to everyone who needs to access the switch, and you must configure and manage password protection on a per-switch basis. In the default configuration, Local authentication is automatically available in the switch. (For more on local authentication, see the password information in the *Configuration and Management Guide* shipped with your switch.
 - **TACACS+ Authentication:** This method enables you to use a TACACS+ server in your network to assign a unique password, user name, and privilege level to each individual or group who needs access to one or more TACACS-aware switches or other TACACS-aware devices. This allows you to administer primary authentication from a central server, and to do so with more options than you have when using only local authentication. (You will still need to use local authentication as a backup if your TACACS+ servers become unavailable.) This means, for example, that you can use a central TACACS+ server to grant, change, or deny access to a specific individual on a specific switch instead of having to change local password assignments on the switch itself, and then have to notify other users of the change.

- **No Security:** The switch can be accessed by anyone without requiring authentication. This is the case when TACACS+ is not enabled on the switch and a local, *manager-level* password is not configured in the switch. Allowing the switch to operate in this mode is not recommended because it compromises switch and network access security.
-

General System Requirements

To use TACACS+ authentication, you need the following:

- Release C.09.*xx* or later software running on your HP ProCurve switch 1600M, 2400M, 2424M, 4000M, or 8000M. Use the following method to view the current software version:

From the Main Menu, click on

1. Status and Counters . . .

1. General System Information

(Check the version number on the **Firmware revision** line.)

Web Browser Interface:

Click on the **Identity** tab.

- A TACACS+ server application installed and configured on one or more servers or management stations in your network. (There are several TACACS+ software packages available.)
- A switch configured for TACACS+ authentication, with access to one or more TACACS+ servers.

Notes

The effectiveness of TACACS+ security depends on correctly using your TACACS+ server application. For this reason HP recommends that you thoroughly test all TACACS+ configurations used in your network.

TACACS+-aware HP switches include the capability of configuring multiple backup TACACS+ servers. HP recommends that you use a TACACS+ server application that supports a redundant backup installation. This allows you to configure the switch to use a backup TACACS+ server if it loses access to the first-choice TACACS+ server.

In release C.09.xx, TACACS+ does not affect web browser interface access. See "Controlling Web Browser Interface Access" on page 28.

General Authentication Setup Procedure

It is important to test the TACACS+ service before fully implementing it. Depending on the process and parameter settings you use to set up and test TACACS+ authentication in your network, you could accidentally lock all users, including yourself, out of access to a switch. While recovery is simple, it may pose an inconvenience that can be avoided. To prevent an unintentional lockout, use a procedure that configures and tests TACACS+ protection for one access type (for example, Telnet access), while keeping the other access type (console, in this case) open in case the Telnet access fails due to a configuration problem. The following procedure outlines a general setup procedure.

Note

If a complete access lockout occurs on the switch as a result of a TACACS+ configuration, see "Troubleshooting TACACS+ Operation" on page 29 for recovery methods.

1. Familiarize yourself with the requirements for configuring your TACACS+ server application to respond to requests from the selected switch. (Refer to the documentation provided with the TACACS+ server software.) This includes knowing whether you need to configure an encryption key. (See "Using the Encryption Key" on page 27.)
2. Determine the following:
 - The IP address(es) of the TACACS+ server(s) you want the switch to use for authentication. If you will use more than one server, determine which server is your first-choice for authentication services.
 - The encryption key(s), if any, for allowing the switch to communicate with the server. You can use either a global key or a per-server key, depending on the encryption configuration in the TACACS+ server(s).
 - The number of log-in attempts you will allow before closing a log-in session. (Default: 3)
 - The period you want the switch to wait for a reply to an authentication request before trying another server.
 - The username/password pairs you want the TACACS+ server to use for controlling access to the switch.
 - The privilege level you want for each username/password pair administered by the TACACS+ server for controlling access to the switch.
 - The passwords you want to use for local authentication (one each for Operator and Manager levels).
3. Plan and enter the TACACS+ server configuration needed to support TACACS+ operation for Telnet access (login/read-only and enable/read-write) to the switch. This includes the username/password sets for logging in at the read-only privilege level and the sets for logging in at the read/write privilege level.

Note on Privilege Levels

When a TACACS+ server authenticates an access request from a switch, it includes a privilege level code for the switch to use in determining which privilege level to grant to the terminal requesting access. The switch interprets a privilege level code of "15" as authorization for the Manager (read/write) privilege level access. Privilege level codes of 14 and lower result in Operator (read-only) access. Thus, when configuring the TACACS+ server response to a request that includes a username/password pair that should have Manager privileges, you must use a privilege level of 15. For more on this topic, refer to the documentation you received with your TACACS+ server application.

If you are a first-time user of the TACACS+ service, HP recommends that you configure only the minimum feature set required by the TACACS+ application to provide service in your network environment. After you have success with the minimum feature set, you may then want to try additional features that the application offers.

4. Ensure that the switch has the correct local password for Manager access. (If the switch cannot find any designated TACACS+ servers, the local manager and operator passwords can be used as the second access control method, depending on the configuration.)

Caution

You should ensure that the switch has a local Manager password. Otherwise, if authentication through a TACACS+ server fails for any reason, then unprotected access will be available through the console port or Telnet.

5. Using a terminal device connected to the switch's console port, configure the switch for TACACS+ authentication *only* for Telnet Login (read-only) access and Telnet Enable access. At this stage, do not configure TACACS+ authentication for console access to the switch, as you may need to use the console for access if the configuration for the Telnet method needs debugging.
6. Ensure that the switch is configured to operate on your network and can communicate with your first-choice TACACS+ server. (At a minimum, this requires IP addressing and a successful ping test from the switch to the server.)
7. On a remote terminal device, use Telnet to attempt to access the switch. If the attempt fails, use the console access to check the TACACS+ configuration on the switch. If you make changes in the switch configuration, check Telnet access again. If Telnet access still fails, check the configuration in your TACACS+ server application for mis-configurations or missing data that could affect the server's interoperation with the switch.

8. After your testing shows that Telnet access using the TACACS+ server is working properly, configure your TACACS+ server application for console access. Then test the console access. If access problems occur, check for and correct any problems in the switch configuration, and then test console access again. If problems persist, check your TACACS+ server application for mis-configurations or missing data that could affect the console access.

Viewing and Configuring TACACS+ Authentication on the Switch

Before You Begin

If you are new to TACACS+ authentication, HP recommends that you read the "General Authentication Setup Procedure" (page 11) and configure your TACACS+ server(s) before configuring authentication on the switch.

Overview

The switch uses three screens for viewing and configuring TACACS+ operation. To access these screens, go to the Authentication Services screen:

From the Main menu, select:

- 2. Switch Management Access Configuration (IP, SNMP, Console) . . .**
- 7. Authentication Services**

```

                                DEFAULT_CONFIG
----- CONSOLE - MANAGER MODE -----
                Switch Configuration - Authentication Services Menu

1. Authentication
2. TACACS+ Operation
3. TACACS+ Servers
4. Return to Previous Menu...
0. Return to Main Menu...

Configures Authentication parameters.
To select menu item, press item number, or highlight item and press <Enter>.

```

Figure 2. The Authentication Services Menu Screen

- **1. Authentication:**
 - Includes the primary and secondary authentication modes (TACACS+ or Local) the switch uses for login (read-only) and enable (read-write) access through the switch's console port and through Telnet access. For definitions of TACACS+ and Local modes, see "Terminology and Options Used in TACACS Applications:" on page 9.
 - Enables 1 to 10 attempts to allow per session (default: 3).
- **2. TACACS+ Operation :**
 - Specifies the timeout (wait) period for a response to an authentication request (default: 5 seconds).

- Defines the (global) encryption key to use if per-server encryption keys are not assigned. (If the TACACS+ servers do not use encryption keys, this parameter should remain empty.) For more on encryption keys, see “Encryption Options in the Switch” on page 27.

■ **3. TACACS+ Servers:**

- Identifies the IP address(es) of the TACACS+ servers to use and the order of priority in which the switch searches for a TACACS+ server.
- Defines the per-server encryption key (if any) associated with each individual server. If a global encryption key is also configured in the TACACS+ Operation screen, the key configured per-server in the TACACS+ Servers screen overrides the global key.

Viewing and Configuring Authentication Options

The Authentication screen displays and configures the access control for console port and Telnet access to the switch. That is, for both access methods, this screen specifies whether to use a TACACS+ server or the switch’s local authentication, or (for some secondary scenarios) no authentication (meaning that if the primary mode fails, authentication is denied). This screen also reconfigures the number of unsuccessful access attempts (incorrect username/password attempts) to allow in a session.

Primary/Secondary Authentication Pairs

For every primary authentication mode there is a compatible secondary authentication mode. Primary authentication can be configured only for Local or TACACS+. Secondary authentication can be configured only for None or Local. To prevent the possibility of a TACACS+ failure locking you out of the switch, there must always be a Local setting in both the Console Login and Console Enable configuration. This means that if there was a problem with the TACACS+ server, you could still access the switch by connecting a terminal device to the Console port and using the switch’s locally configured passwords. Table 1, below, shows the options for authentication pairs.

Table 1. Primary/Secondary Authentication Table

Access Method and Privilege Level	Authentication Options		Effect on Access Attempts
	Primary	Secondary	
Console — Login	local	none	Local username/password access only.
	tacacs	local	If Tacacs+ server unavailable, uses local username/password access.
Console — Enable	local	none	Local username/password access only.
	tacacs	local	If Tacacs+ server unavailable, uses local username/password access.
Telnet — Login	local	none	Local username/password access only.
	tacacs	local	If Tacacs+ server unavailable, uses local username/password access.
	tacacs	none	If Tacacs+ server unavailable, denies access.
Telnet — Enable	local	none	Local username/password access only.
	tacacs	local	If Tacacs+ server unavailable, uses local username/password access.
	tacacs	none	If Tacacs+ server unavailable, denies access.

Caution Regarding the Use of Local for Login Primary Access

During local authentication (which uses passwords configured in the switch instead of in a TACACS+ server), the switch grants read-only access if you enter the Operator password, and read-write access if you enter the Manager password. For example, if you configure authentication on the switch with Telnet Login Primary as **Local** and Telnet Enable Primary as **Tacacs**, when you attempt to Telnet to the switch, you will be prompted for a local password. If you enter the switch's local Manager password (or, if there is no local Manager password configured in the switch) you can bypass the TACACS+ server authentication for Telnet Enable Primary and go directly to read-write (Manager) access. Thus, for either the Telnet or console access method, configuring Login Primary for Local authentication while configuring Enable Primary for TACACS+ authentication is not recommended, as it defeats the purpose of using the TACACS+ authentication. If you want Enable Primary log-in attempts to go to a TACACS+ server, then you should configure both Login Primary and Enable Primary for Tacacs authentication instead of configuring Login Primary to Local authentication.

To View or Configure the Authentication Parameters for either Console or Telnet Access

This procedure displays and configures the switch's authentication modes and the number of log-in attempts to allow per session.

1. From the Main menu, select:
 2. **Switch Management Access Configuration (IP, SNMP, Console) . . .**
 7. **Authentication Services**
 1. **Authentication**

```

                                DEFAULT_CONFIG
-----
Switch Configuration - Authentication Services - Authentication
-----
Login Attempts : 3

Task Name   Login Primary   Login Secondary   Enable Primary   Enable Secondary
----- + -----
Console    | Local          None              Local            None
Telnet     | Local          None              Local            None

Actions->  Cancel   Edit   Save   Help

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

Figure 3. The Default Authentication Screen

Note

As described under "General Authentication Setup Procedure" on page 11, HP recommends that you configure, test, and troubleshoot authentication via Telnet access before you configure authentication via console port access. This helps to prevent accidentally locking yourself out of switch access due to errors or problems in setting up authentication in either the switch or your TACACS+ server.

2. Press [E] (for **Edit**) to move the cursor to the **Login Attempts** field (default: **3**; range **1 - 10**). To change this value, type a new number.

3. Use the downarrow key to select the **Login Primary** field for the access method you are configuring (Console or Telnet).
4. Use the Space bar to select the **Login Primary** authentication mode (TACACS or Local).

Note

For a particular access method (Console or Telnet), if you want to use the TACACS mode for **Enable Primary**, then HP recommends that you configure both the **Login Primary** and the **Enable Primary** as **Tacacs**. See "Caution Regarding the Use of Local for Login Primary Access" on page 16.

5. Use the rightarrow key to move the cursor to the **Login Secondary** column for the selected access method (Console or Telnet), then use the Space bar to select the secondary authentication mode.
6. Use the rightarrow key to move the cursor to the **Enable Primary** column, then use the Space bar to select the **Enable Primary** authentication mode (TACACS or Local).
7. Use the rightarrow key to move the cursor to the **Enable Secondary** column for the selected access method (Console or Telnet), then use the Space bar to select the secondary authentication mode.

For example, if you use steps 2 through 7 to configure **Login Attempts** for 2 and Telnet access with **Tacacs** and **None** for authentication modes in both Login and Enable access levels, the Authentication screen would appear as follows:

```

                                DEFAULT_CONFIG
-----
Switch Configuration - Authentication Services - Authentication

Login Attempts : 2

Task Name   Login Primary   Login Secondary   Enable Primary   Enable Secondary
-----+-----
Console    | Local           None              Local            None
Telnet     | Tacacs          None              Tacacs           None

Actions->   _ancel         Edit         _ave           _elp

Edit the fields displayed above.
Use arrow keys to change action selection and <Enter> to execute action.

```

Figure 4. Example of Authentication Configuration for Telnet Access

8. Press [S] (for **S**ave) to save your configuration changes.

9. Do the following:

- a. If you need to view or change the period the switch waits for a TACACS+ server to respond to an authentication request, or if you need to view or configure a global encryption key, go to “Viewing and Configuring the Timeout Period and (Global) Encryption Key” on page 19. (For more on encryption keys, see “Encryption Options in the Switch” on page 27.)
- b. To view or configure the IP address(es) of the specific TACACS+ server(s) the switch is using for authentication services, or to view or configure server-specific encryption keys, go to “Viewing and Configuring TACACS+ Servers and (Per-Server) Encryption Keys” on page 20. (For more on encryption keys, see “Encryption Options in the Switch” on page 27.)

Viewing and Configuring the Timeout Period and (Global) Encryption Key

Timeout Period. After polling a TACACS+ server, the switch uses a configurable timeout period to determine how long to wait for a response. (The default is five seconds.) If no response is received during the timeout period, the switch tries again using a different TACACS+ server (if configured in the TACACS+ Servers screen).

(Global) Encryption Key. When configured, the switch uses this optional key whenever you have not configured a per-server key in the TACACS+ Servers screen (page 20). (For more on encryption keys, see “Encryption Options in the Switch” on page 27.)

Note

Use a global encryption key only if all TACACS+ servers supporting authentication for a particular switch have keys configured for that switch. If a TACACS+ server supporting the switch does not have an encryption key applicable to the switch, then assigning a global key in the switch will block authentication support from that server.

Viewing and Configuring the Timeout Period and (Global) Encryption Key

To View or Configure the Timeout period and global encryption key:

1. From the Main menu, select:
 2. **Switch Management Access Configuration (IP, SNMP, Console) ...**
 7. **Authentication Services**
 2. **TACACS+ Operation**




Figure 5. The Default Screen for Changing Timeout and/or Configuring the Optional Global Encryption Key

2. To change the TACACS+ Operation parameters:
 - a. Press [E] (for **Edit**). The cursor moves to the **Timeout** field.
 - b. Type the timeout value you want (any value from 1 to 255 seconds).
3. To enter or change the (optional) global encryption key
 - a. Press the downarrow key to move the cursor to the **Encryption Key** field.
 - b. Type the global encryption key.
4. Press [Enter] to return to the Actions bar, then [S] (for **Save**) to save your configuration changes and return to the Authentication Services Menu.

Viewing and Configuring TACACS+ Servers and (Per-Server) Encryption Keys

The TACACS+ Servers screen lists up to three IP addresses of the TACACS+ servers you want the switch to use for authentication, in order of priority, along with any per-server encryption keys.

Using (Per-Server) Encryption Keys

If you are assigning the switch to a TACACS+ server that uses an encryption key, you should enter that key next to the server's IP address if either of the following is true:

- You are not assigning a Global encryption key in the TACACS+ Operation screen (figure 5 on page 20).
- You are assigning a global encryption key in the TACACS+ Operation screen, but the global key is not the same as the key used by the server you are currently entering.

For more on encryption keys, see “Encryption Options in the Switch” on page 27.

Viewing and Configuring TACACS+ Server IP Addresses and (Per-Server) Encryption Keys

To View or Configure the Timeout period and global encryption key:

1. From the Main menu, select:
 2. **Switch Management Access Configuration (IP, SNMP, Console) . . .**
 7. **Authentication Services**
 2. **TACACS+ Servers**

```

                                DEFAULT_CONFIG
=====-- CONSOLE - MANAGER MODE -----=====
Switch Configuration - Authentication Services - TACACS+ Servers

Server IP Address                Encryption Key
-----
10.28.227.115  secret-1
10.28.227.101

Actions->  Back  Add  Edit  Delete  Help

Add a new record.
Use up/down arrow keys to change record selection, left/right arrow keys to
change action selection, and <Enter> to execute action.
```

This encryption key (secret-1) is for use only with TACACS+ server 10.28.227.115.

Since a per-server key is not assigned to this server, the switch will use the global key for authentication requests to this server (if the key is configured in the TACACS+ Operation screen). If there is no global key, then authentication requests from the switch to this server will not include an encryption key.

Figure 6. Example of a Prioritized List of TACACS+ Servers

2. To add a server and, if necessary, an encryption key for requests to that server:

- a. Press [A] (for **Add**) to display this screen.

```
                                DEFAULT_CONFIG                                :
-----
----- CONSOLE - MANAGER MODE -----
Switch Configuration - Authentication Services - TACACS+ Servers

Server IP Address : ██████████
Encryption Key :

Actions->  Cancel      Edit      Save      Help
████████████████████████████████████████████████████████████████████████████████
Use arrow keys to change field selection, <Space> to toggle field choices,
and <Enter> to go to Actions.
```

Figure 7. Screen for Adding TACACS+ Servers and Associated (Per-Server) Encryption Keys

- b. Enter the desired IP address in the **Server IP Address** field.
 - c. If you want to configure an encryption key to use with authentication requests to this server, press the downarrow key and type the key in the **Encryption Key** field. Otherwise, skip this step.
 - d. Press [Enter], then [S] (for **Save**) to save the new server IP and optional key in the switch's TACACS+ server configuration and return to the TACACS+ Server screen (figure 6).
3. To delete a server entry:
 - a. Go to the TACACS+ Server screen (figure 6) and highlight the IP address of the server you want to delete.
 - b. Press [D] (for **Delete**). The screen then prompts you as shown below:

```

                                DEFAULT_CONFIG
----- -- CONSOLE - MANAGER MODE -----
Switch Configuration - Authentication Services - TACACS+ Servers

Server IP Address                Encryption Key
----- --
10.28.227.115                    secret-1
10.28.227.108                    secret-2
10.28.227.101

Continue Deletion of record ? No

```

The switch prompts you to verify deletion of the selected TACACS+ server and (if configured) the associated per-switch encryption key. To complete the deletion, use the Space bar to select **Yes**, then press [Enter] to complete the deletion and return to the TACACS+ Server screen.

Use up/down arrow keys to change record selection, left/right arrow keys to change action selection, and <Enter> to execute action.

Figure 8. Example of Deleting a TACACS+ Server and Encryption Key

- c. Use the Space bar to select Yes.
- d. Press [Enter] to complete the deletion and return to the TACACS+ Server screen

Press [B] (for **Back**) to return to the Authentication Services Menu.

Controlling Server Priority when Adding and Removing Server Addresses

This section describes how adding and deleting servers affects how the switch prioritizes TACACS+ servers for authentication requests.

- When the server list contains multiple servers, the switch always tries to authenticate TACACS+ requests through the first server on the list. If the first server does not respond, then the switch tries the second server, and so-on. If none of the servers in the list respond, then the switch tries to authenticate access through local passwords on the switch (if secondary authentication is configured as **Local** for the access method being used).
- When the TACACS+ Server list contains only one or two IP addresses, adding a third IP address places that address at the end of the list. When the list is full (three IP addresses), you must delete one address before you can add another.
- If you delete an IP address from the list and then add a new IP address to the list, the switch gives the new address the same priority as the deleted address. For example, if you have the following list configured:

1.1.1.1
2.2.2.2

If you delete 1.1.1.1 and add 3.3.3.3, then 3.3.3.3 will have the highest priority and the list will appear as follows:

3.3.3.3
2.2.2.2

- If you delete multiple addresses and then add new addresses, the switch assigns priority as follows:
 - a. The first new address receives the priority that belonged to the first of the deleted addresses.
 - b. The second new address receives the priority that belonged to the second of the deleted addresses, and so on.

For example, suppose you have three addresses in the Server IP Address list:

3.3.3.3	First priority.
1.1.1.1	Second priority.
2.2.2.2	Third priority.

In this case, 3.3.3.3 has first priority; 1.1.1.1 has second priority, and so on. If you delete 1.1.1.1 from the list, you will then have:

3.3.3.3	First priority.
2.2.2.2	Second priority.

If you then add a new IP address, 4.4.4.4, this new address would assume second priority in the list:

3.3.3.3	First priority.
4.4.4.4	Second priority.
2.2.2.2	Third priority.

Changing Server IP Address Priority While Keeping the Same List of Servers

1. Delete the server you want to move down in priority.
2. Delete the server you want to move up in priority.
3. Add the server you deleted in step 2.
4. Add the server you deleted in step 1.

How Authentication Operates

General Authentication Process Using a TACACS+ Server

Authentication through a TACACS+ server operates generally as described below. For specific operating details, refer to the documentation you received with your TACACS+ server application.

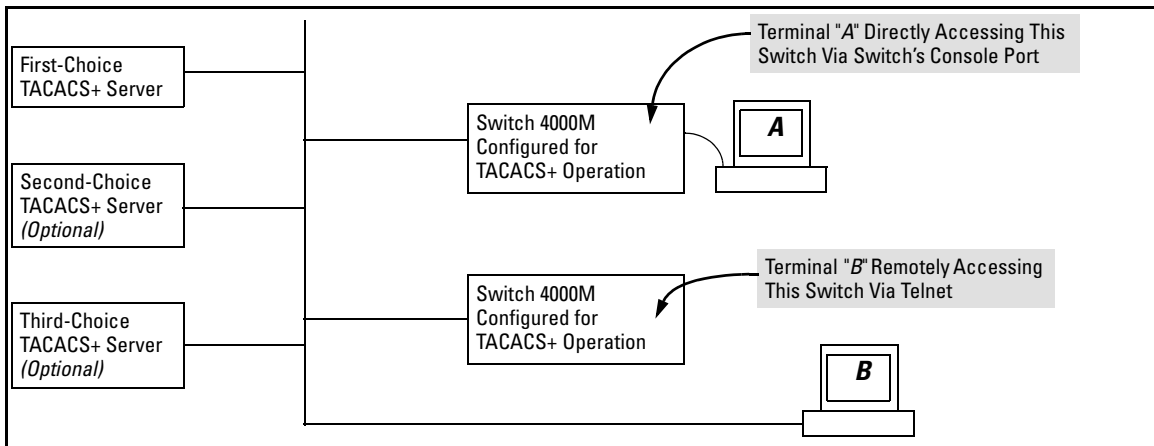


Figure 9. Using a TACACS+ Server for Authentication

Using figure 9, above, after either switch detects a logon request from a remote or directly connected terminal, the following events occur:

1. The switch queries the first-choice TACACS+ server for authentication of the request.
 - If the switch does not receive a response from the first-choice TACACS+ server, it attempts to query a secondary server. If the switch does not receive a response from any TACACS+ server, then it uses its own local username/password pairs to authenticate the logon request. (See "Local Authentication Process", on page 26.)
 - If a TACACS+ server recognizes the switch, it forwards a username prompt to the requesting terminal via the switch.
2. When the requesting terminal responds to the prompt with a username, the switch forwards it to the TACACS+ server.
3. After the server receives the username input, the server forwards a password prompt to the requesting terminal via the switch.

4. When the requesting terminal responds to the prompt with a password, the switch forwards it to the TACACS+ server and one of the following actions occurs:
 - If the username/password pair received from the requesting terminal matches a username/password pair previously stored in the server, then the server passes access permission through the switch to the terminal.
 - If the username/password pair entered at the requesting terminal does not match a username/password pair previously stored in the server, access is denied. In this case, the terminal is again prompted to enter a username and repeat steps 2 through 4. In the default configuration, the switch allows up to three attempts to authenticate a login session. If the requesting terminal exhausts the attempt limit without a successful TACACS+ authentication, the login session is terminated and the operator at the requesting terminal must initiate a new session before trying again.

Local Authentication Process

The switch uses local authentication only if one of these two conditions exists:

- "Local" is the primary authentication option for the access method being used.
- TACACS+ is the primary authentication mode for the access method being used. However, the switch was unable to connect to any TACACS+ servers (or no servers were configured) *AND Local* is the secondary authentication mode being used.

(For a listing of authentication options, see table 1 on page 16.)

For local authentication, the switch uses the operator-level and manager-level passwords previously configured locally on the switch. (These are the passwords you can configure using the menu interface or the web browser interface—which enables only the local password configuration).

- If the operator at the requesting terminal correctly enters the password for either access level, access is granted.
- If the password entered at the requesting terminal does not match either password previously configured locally in the switch, access is denied. In this case, the terminal is again prompted to enter a password. In the default configuration, the switch allows up to three attempts. If the requesting terminal exhausts the attempt limit without a successful authentication, the login session is terminated and the user at the requesting terminal must initiate a new session before trying again.

Note

The switch's menu allows you to configure only the local Operator and Manager passwords. In this case, all prompts for local authentication will request only a local password. However, if you use the web browser interface to configure usernames for local access, you will be prompted for both a local username and a local password during local authentication.

TACACS+ does not affect web browser interface access. To block unauthorized access through the web browser interface, see "Controlling Web Browser Interface Access When Using TACACS+ Authentication" on page 28.

Using the Encryption Key

General Operation

When used, the encryption key (sometimes termed "key", "secret key", or "secret") helps to prevent unauthorized intruders on the network from reading username and password information in TACACS+ packets moving between the switch and a TACACS+ server. At the TACACS+ server, a key may include both of the following:

- **Global key:** A general key assignment in the TACACS+ server application that applies to all TACACS-aware devices for which an individual key has not been configured.
- **Individual key:** A unique key assignment in the TACACS+ server application that applies to a specific TACACS-aware device.

Note

Configure a key in the switch only if the TACACS+ server application has this exact same key configured for the switch. That is, if the key parameter in switch "X" does not exactly match the key setting for switch "X" in the TACACS+ server application, then communication between the switch and the TACACS+ server will fail.

Thus, on the TACACS+ server side, you have a choice as to how to implement a key. On the switch side, it is necessary only to enter the key parameter so that it exactly matches its counterpart in the server. For information on how to configure a general or individual key in the TACACS+ server, refer to the documentation you received with the application.

Encryption Options in the Switch

If a TACACS+ server supporting a particular switch is configured to use an encryption key when authenticating access to the switch, then you must include the key in the switch's authentication configuration. There are two options for configuring encryption keys in the switch: Global and per-server. If all TACACS+ servers supporting the switch apply the same encryption key when authenticating access to the switch, then you can specify one, global key. (See figure 5 on page 20.) In the case of multiple TACACS+ servers supporting the switch with different encryption keys, then you must configure the keys on a per-server basis. (See figures 6 and 7 on page 21 and 22.)

Note

If you configure both a global key and one or more per-server keys, the per-server keys will override the global key for the specified servers.

Controlling Web Browser Interface Access When Using TACACS+ Authentication

In release C.09.*xx*, configuring the switch for TACACS+ authentication does not affect web browser interface access. To prevent unauthorized access through the web browser interface, do one or more of the following:

- Configure local authentication (a Manager user name and password and, optionally, an Operator user name and password) on the switch.
- Configure the switch's Authorized IP Manager feature to allow web browser access only from authorized management stations. (The Authorized IP Manager feature does not interfere with TACACS+ operation.)
- Disable web browser access to the switch by going to the System Information screen in the Menu interface and configuring the **Web Agent Enabled** parameter to **No**.

Operating Notes

- If you configure Authorized IP Managers on the switch, it is not necessary to include any devices used as TACACS+ servers in the authorized manager list. That is, TACACS+ operates regardless of any Authorized IP Manager configuration.
- When the switch is not configured to use TACACS+ servers—or when the switch's only designated TACACS+ servers are not accessible—*setting a local Operator password without also setting a local Manager password does not protect the switch from manager-level access by unauthorized persons.* See also the **Caution** on page 16.

Troubleshooting TACACS+ Operation

Event Log. When troubleshooting TACACS+ operation, check the switch's Event Log (accessed from the Main menu) for indications of problem areas.

All Users Are Locked Out of Access to the Switch. If the switch is functioning properly, but no username/password pairs result in console or Telnet access to the switch, the problem may be due to how the TACACS+ server and/or the switch are configured. Use one of the following methods to recover:

- Access the TACACS+ server application and adjust or remove the configuration parameters controlling access to the switch.
- Disconnect the switch from network access to any TACACS+ servers and then log in to the switch using either Telnet or direct console port access. Because the switch cannot access a TACACS+ server, it will default to local authentication. You can then use the switch's local Operator or Manager password to log on.
- As a last resort, use the Clear/Reset button combination to reset the switch to its factory default boot-up configuration. Taking this step means you will have to reconfigure the switch to return it to operation in your network.

No Communication Between the Switch and the TACACS+ Server Application. If the switch can access the server device (that is, it can **ping** the server), then a configuration error may be the problem. Some possibilities include:

- The server IP address configured in the switch's TACACS+ Server screen may not be correct.
- The encryption key configured in the server does not match the encryption key configured in the switch. Verify the key in the server and compare it to the key configured in the switch.
- The accessible TACACS+ servers are not configured to provide service to the switch.

Access Is Denied Even Though the Username/Password Pair Is Correct. Some reasons for denial include the following parameters controlled by your TACACS+ server application:

- The account has expired.
- The access attempt is through a port that is not allowed for the account.
- The time quota for the account has been exhausted.
- The time credit for the account has expired.
- The access attempt is outside of the timeframe allowed for the account.
- The allowed number of concurrent logins for the account has been exceeded.

For more help, refer to the documentation provided with your TACACS+ server application.

Unknown Users Allowed to Login to the Switch. Your TACACS+ application may be configured to allow access to unknown users by assigning them the privileges included in a *default user* profile. Refer to the documentation provided with your TACACS+ server application.

System Allows Fewer Login Attempts than Specified in the Switch Configuration. Your TACACS+ server application may be configured to allow fewer login attempts than you have configured in the switch.

CDP

Introduction

In HP ProCurve switches running software version C.09.*xx* or later, CDP-v1 (Cisco Discovery Protocol, version 1) provides data that aids SNMP-based network mapping utilities designed to discover devices running CDP in a network. To make this data available, the switch transmits information about itself via CDP packets to adjacent devices, and also receives and stores information about adjacent devices running CDP. This enables each CDP device to receive and maintain identity data on each of its CDP neighbors and pass this information off to an SNMP utility designed to query the CDP area of the device's MIB.

Note

To take advantage of CDP, you should have a working knowledge of SNMP operation and an SNMP utility capable of polling the switches for CDP data. HP's implementation of CDP places specific data into the switch's Management Information Base (MIB). However, retrieval of this data for network mapping is dependent on the operation of your SNMP utility. Refer to the documentation provided with the utility.

An SNMP utility can progressively discover CDP devices in a network by:

1. Reading a given device's CDP Neighbor table (in the Management Information Base, or MIB) to learn about other, neighbor CDP devices
2. Using the information learned in step 1 to go to and read the neighbor devices' CDP Neighbors tables to learn about additional CDP devices, and so on

This section describes CDP operation in the HP ProCurve Switches 1600M, 2400M, 2424M, 4000M, and 8000M. For information on how to use an SNMP utility to retrieve the CDP information from the switch's CDP Neighbors table (in the switch's MIB), refer to the documentation provided with the particular SNMP utility. For information on the object identifiers in the CDP MIB, see "Viewing the CDP Neighbor Table" on page 39.

CDP Terminology

- **CDP Device:** A switch, server, router, workstation, or other device running CDP.
- **CDP-Aware:** A device that has CDP in its operating code (with CDP either enabled or disabled in that device).
- **CDP-Disabled:** A CDP-aware device on which CDP is currently disabled.

- **Non-CDP Device:** A device that does not have CDP capability in its operating code.
- **CDP Neighbor:** A CDP device that is either directly connected to another CDP device or connected to that device by a non-CDP device, such as some hubs.

General CDP Operation

The switch stores information about adjacent CDP devices in a *CDP Neighbors table* maintained in the switch's MIB (Management Information Base). This data is available to SNMP-based applications designed to read CDP data from the MIB. For example:

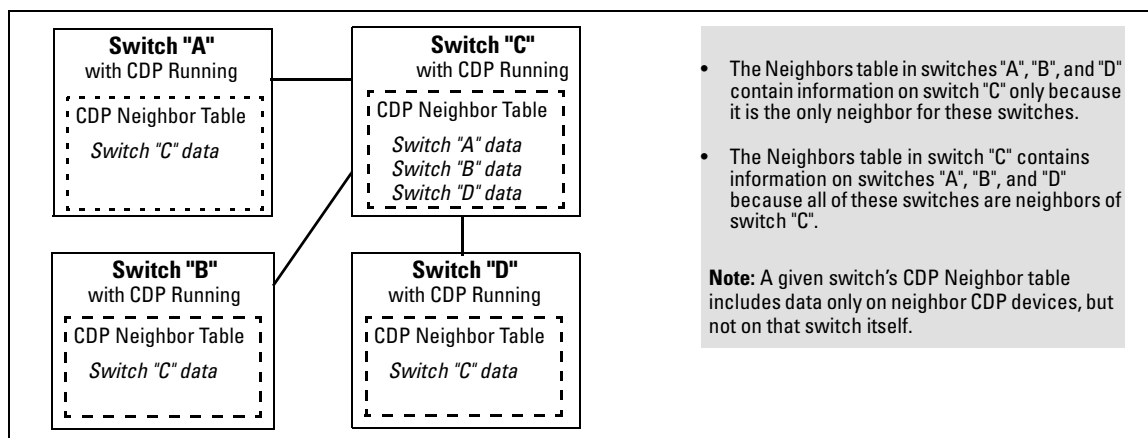


Figure 10. Example of How the Switches Store Data on Neighbor CDP Devices

Outgoing Packets

A switch running CDP periodically transmits a one-hop CDP packet out each of its ports. This packet contains data describing the switch and, if the one-hop destination is another device running CDP, the receiving device stores the sending device's data in a CDP Neighbors table. The receiving device also transmits a similar one-hop CDP packet out each of its ports to make itself known to other CDP devices to which it is connected. Thus, each CDP device in the network provides data on itself to the CDP neighbors to which it is directly connected. However, there are instances where a packet is forwarded beyond the immediate neighbor, or simply dropped.

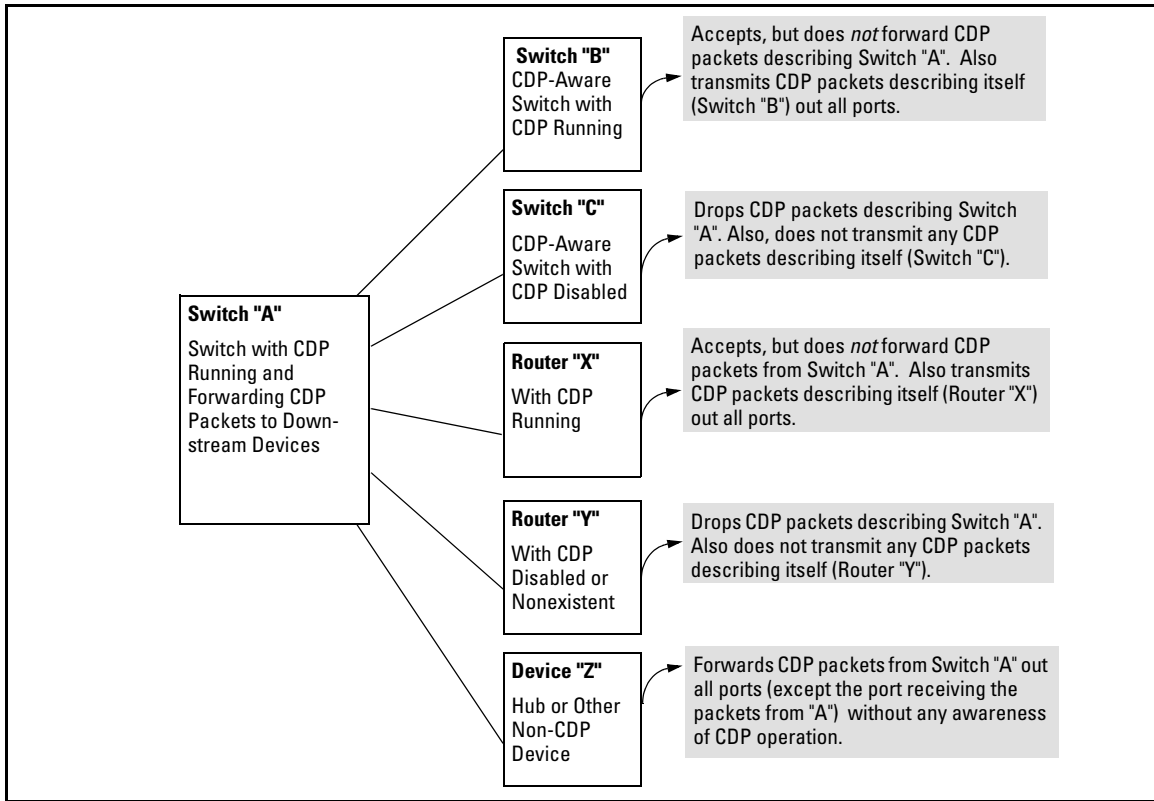


Figure 11. Example of Outgoing CDP Packet Operation

Incoming CDP Packets

When a CDP-enabled switch receives a CDP packet from another CDP device, it enters that device's data in the CDP Neighbors table, along with the port number where the data was received (and does not forward the packet). The switch also periodically purges the table of any entries that have expired. (The hold time for any data entry in the switch's CDP Neighbors table is configured in the device transmitting the CDP packet, and cannot be controlled in the switch receiving the packet.) The HP ProCurve switches review the list of CDP neighbor entries every three seconds, and purge any expired entries.

Non-CDP devices such as some hubs and other devices that do not have CDP capability are transparent to CDP operation. (Other hubs are CDP-aware, but still forward CDP packets as if they were transparent to CDP operation. See "CDP-Capable Hubs" on page 41.) However, an intervening CDP-aware device that is CDP-disabled is *not* transparent. For example, in figure 12, the CDP neighbor pairs are as follows: A/1, A/2, A/3, A/B, B/C. Note that "B" and "E" are *not* neighbors because the intervening CDP-disabled switch "D" does not forward CDP packets; i.e. is not transparent to CDP traffic. (For the same reason, switch "E" does not have any CDP neighbors.)

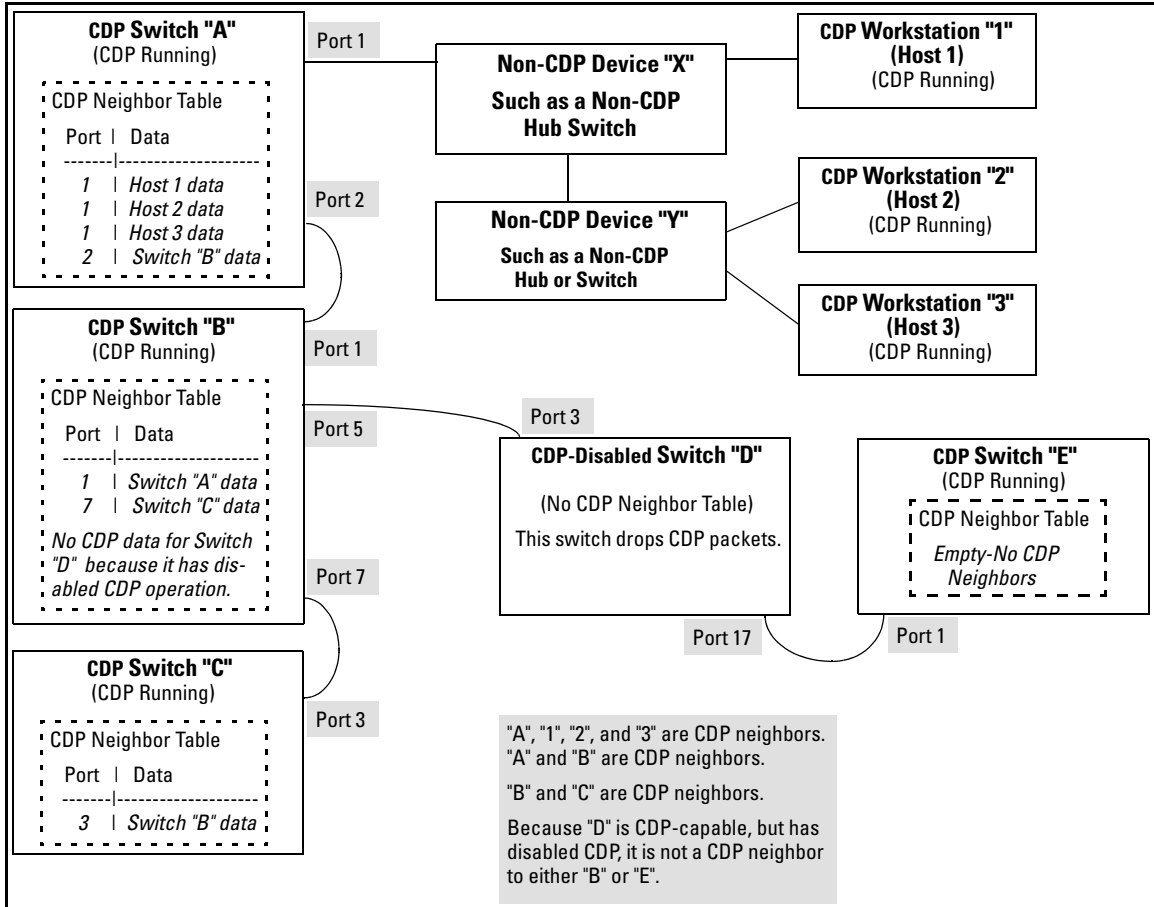


Figure 12. Example of Incoming CDP Packet Results

Thus, based on the CDP packets it receives, each CDP device maintains a per-port data entry for each of its neighbors that are running CDP, but not for other CDP devices that are accessible only through a CDP neighbor. (See the relationship between switches A, B, and C in figure 12.) In other words, a CDP device will have data on its immediate CDP neighbors (including those reached through a device that is transparent to CDP), but not to other CDP devices in the network.

Table 2. How Devices Handle Incoming CDP Packets

Status of Device Receiving a CDP Packet	Action of Receiving Device
Running CDP	Stores neighbor data in CDP Neighbor table in the device MIB. Does not forward CDP packet.
CDP Disabled	Drops CDP packet. There is no CDP Neighbor table in the device MIB and no CDP neighbor data is stored.
No CDP Capability	Forwards CDP packet out all ports except the port on which the packet was received.
Router Running CDP	Stores neighbor data in the CDP Neighbor table in the router's MIB. Does not forward CDP packet.
Router with CDP (1) Disabled or (2) Not CDP-Capable	Drops CDP packet.

Non-CDP devices (that is, devices that are not capable of running CDP) are transparent to CDP operation. However, an intervening CDP-aware device that is CDP-disabled is *not* transparent. For example, in figure 12 (page 34), "B", "D", and "E" are *not* CDP neighbors because "D" (the intervening CDP-disabled switch) does not forward CDP packets; i.e. is not transparent to CDP traffic. (For the same reason, switch "E" does not have any CDP neighbors.)

Figure 12 (page 34) illustrates how multiple CDP neighbors can appear on a single port. In this case, switch "A" has three CDP neighbors on port 1 because the intervening devices are not CDP-capable and simply forward CDP neighbors data out all ports (except the port on which the data was received).

Configuring CDP on the Switch

On the HP ProCurve switches 1600M, 2400M, 2424M, 4000M, and 8000M you can:

- View the switch's current CDP configuration.
- Enable or disable CDP (Default: Enabled).
- Specify the hold time (CDP packet time-to-live) for CDP data delivered to neighboring CDP devices. For example, in CDP switch "A" you can specify the hold time for switch "A" entries in the CDP Neighbor tables of other CDP devices. (Default: 180 seconds)
- Specify the transmission interval for CDP packets. (Default: 60 seconds).
- Use the **walkmib** command to display the current contents of the switch's CDP Neighbors table (page 39).

Viewing and Changing the Switch's Current CDP Configuration

Parameter	Operation
Enable CDP	<p>Enabling CDP operation (the default) on the switch causes the switch to:</p> <ul style="list-style-type: none"> • Transmit CDP packets describing itself to other, neighboring CDP devices • Add entries to its CDP Neighbors table for any CDP packets it receives from other, neighboring CDP devices <p>Disabling CDP operation clears the switch's CDP Neighbors table, prevents the switch from transmitting outbound CDP packets to advertise itself to neighboring CDP devices, and causes the switch to drop inbound CDP packets from other devices without entering the data in the CDP Neighbors table in the switch's MIB (Management Information Base).</p>
CDP Hold Time	The default hold time for the switch's CDP packet information in the CDP Neighbors table of another CDP device is 180 seconds (range: 10 - 255 seconds). This parameter is controlled in the transmitting switch, and applies to all outbound CDP packets the switch transmits.
CDP Transmit Interval	The default interval the switch uses to transmit CDP packets describing itself to other, neighbor devices is 60 seconds. (range: 5 - 254).

To view and change the CDP configuration:

1. From the Main Menu, select:

3. Switch Configuration ...

1. System Information

```

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

System Name : DEFAULT_CONFIG          Enable CDP [Yes] : Yes
System Contact :                      CDP Hold Time [180] : 180
System Location :                      CDP Transmit Interval [60] : 60
                                CDP Parameters
Inactivity Timeout (min) [0] : 0      MAC Age Interval (sec) [300] : 300
Web Agent Enabled [Yes] : Yes

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

Actions->  Cancel  Edit  Save  Help

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

```

Figure 13. The Default System Information Screen

2. To change the CDP configuration, press [E] (for **E**dit), and then use the the downarrow key to move to the CDP parameter you want to change.
 - To enable or disable CDP operation, use the Space bar to select **Yes** or **No**.
 - To change the **CDP Hold Time** or **CDP Transmit Interval**, type the desired time value in the appropriate field.
3. Press [Enter] to return to the **Actions** line, then press [S] (for **S**ave) to save your changes.

Effect of Spanning Tree (STP) On CDP Packet Transmission

If STP has blocked a port on the switch, that port does not transmit CDP packets. However, the port still receives CDP packets if the device on the other end of the link has CDP enabled. Thus, for example, if switch "X" has two ports linked to switch "Y" (which is a CDP neighbor and also the STP root device) and STP blocks traffic on one port and forwards traffic on the other:




Figure 14. Example of How STP and the STP Root Device Affects CDP Packet Transmission

- Switch "X" sends outbound CDP packets on the forwarding link, and the switch "Y" CDP Neighbors table records switch "X" on only one port.
- Switch "Y" sends outbound CDP packets on both links, and the switch "X" CDP Neighbors table shows switch "Y" on both ports.

To summarize, in a CDP neighbor pair running STP with redundant links, if one of the switches is the STP root, it transmits CDP packets out all ports connecting the two switches, while the other switch transmits CDP packets out only the unblocked port. Thus, the STP root switch will appear on multiple ports in the non-root switch's CDP Neighbors table, while the non-root switch will appear on only one port in the root switch's CDP Neighbors table.

How the Switch Selects the IP Address To Include in Outbound CDP Packets

A switch with CDP enabled uses the following prioritized criteria to determine which IP address to include in its outbound CDP packets:

1. If VLAN support is disabled (the default setting) and the switch is configured with an IP address, then the outbound CDP packets from any port on the switch will carry that IP address.
2. If VLANs are enabled and the default VLAN has an IP address, then the outbound CDP packets from any port belonging to the default VLAN will carry the IP address of the default VLAN.

Note: If a switch has the capability to designate a VLAN other than the default VLAN as the Primary VLAN (such as the HP ProCurve Series 2500 switches), then item 2, above, applies to the switch's Primary VLAN.

3. If a port belongs to only one VLAN that has an IP address, then the port includes that address in the outbound CDP packets from that port.
4. If 1, 2, or 3 do not apply to a port, and the port belongs to multiple VLANs having IP addresses, then outbound CDP packets from that port will carry the IP address of the VLAN with the lowest VID (VLAN Identification number). (See figure 15, below.)
5. If a port does not belong to any VLANs that have an IP address, then outbound CDP packets from that port carry the loopback IP address (127.0.0.1).

For example, in figure 15, port 1 on CDP switch "X" is connected to port 5 on CDP neighbor switch "Y", with the indicated VLAN configuration on port 5:

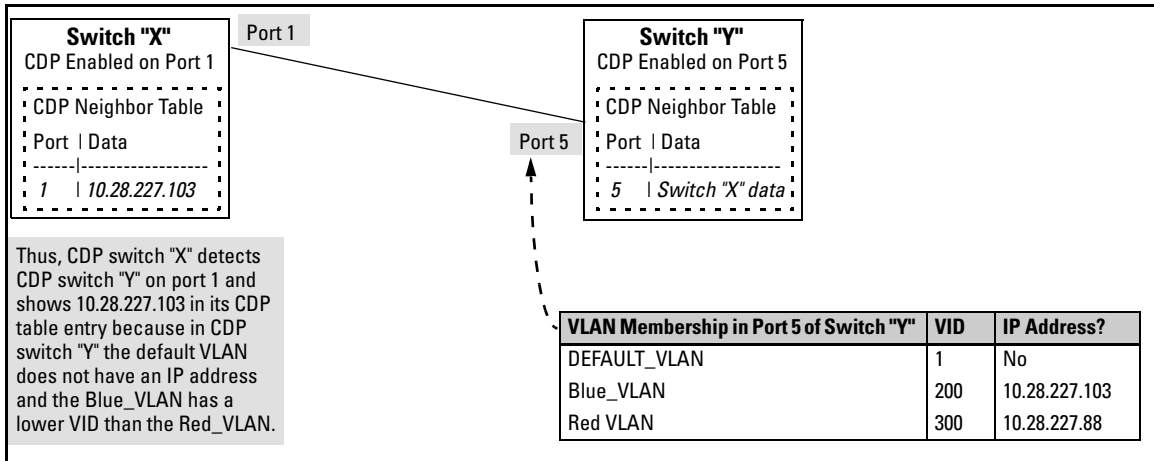


Figure 15. Example of IP Address Selection when the CDP Neighbor Has Multiple VLANs with IP Addresses

Viewing the CDP Neighbor Table

The switch places the data received from inbound CDP packets into its MIB (Management Information Base). To display this information, use the **walkmib** command.

Table 3. CDP Neighbors Data

CDP Neighbor Data	MIB Value
Address Type	Always "1" (IP address only). (MIB: cdpCacheAddressType.<inbound port #>)
CDP Cache Address	IP address of source device in hexadecimal format. (MIB: cdpCacheAddress.<inbound port #>)
Software Version	ASCII String (MIB: cdpCacheVersion.<inbound port #>)
Device Name (ASCII string) and Device MAC Address	In HP ProCurve switches, this is the value configured for the System Name parameter. (MIB: cdpCacheDeviceId.<inbound port #>)
Source Port Number	On the source (neighbor) device, the number of the port through which the CDP packet was sent. (MIB: cdpCacheDevicePort.<inbound port #>)
Product Name (ASCII string)	Platform name designated by vendor. (MIB: cdpCachePlatform.<inbound port #>)
Capability Code (Device Type)	1: Router 2: Transparent Bridge 4: Source Route Bridge 8: Switch 16: Host 32: IGMP conditional filtering 64: Repeater (MIB: cdpCacheCapabilities.<inbound port #>)

Displaying CDP Neighbor Data. Go to the switch's command prompt and use the **walkmib** command, as shown below.

1. From the Main Menu, select:

5. Diagnostics . . .

4. Command Prompt

2. Enter the following walkmib command:

```
walkmib cdpCacheEntry
```

For example, executing the above command in a Switch 8000M connected to two HP ProCurve Series 2500 switches with CDP enabled produces a listing similar to that shown in figure 16:

DEFAULT_CONFIG: walkmib cdpCacheEntry

```
                                DEFAULT_CONFIG
----- CONSOLE - MANAGER MODE -----
cdpCacheAddressType.4.2 = 1
cdpCacheAddressType.7.1 = 1
cdpCacheAddress.4.2 = 7f 00 00 01
cdpCacheAddress.7.1 = 7f 00 00 01
cdpCacheVersion.4.2 = Revision F.02.02 /sw/code/build/info(f00)
cdpCacheVersion.7.1 = Revision F.02.D1 /sw/code/build/info(f00)
cdpCacheDeviceId.4.2 = HP ProCurve Switch 2524(000010-ff00
cdpCacheDeviceId.7.1 = HP ProCurve Switch 2512(0000b8-0000
cdpCacheDevicePort.4.2 = 24
cdpCacheDevicePort.7.1 = 12
cdpCachePlatform.4.2 = HP J4813A ProCurve Switch 2524
cdpCachePlatform.7.1 = HP J4812A ProCurve Switch 2512
cdpCacheCapabilities.4.2 = 8
cdpCacheCapabilities.7.1 = 8
DEFAULT_CONFIG: _
```

The first number after the MIB string is the HP 8000M port on which the data point for that entry was received.

Figure 16. Example of CDP Neighbor Data in the Switch 8000M

For the Switch 1600M/2400M/2424M/4000M/8000M MIB, go to <http://www.hp.com/go/hpprocurve> and click on **software**.

Operating Notes

Neighbor Maximum. The HP ProCurve switches support up to 60 neighbors in the CDP Neighbors table.

Multiple CDP Devices on the Same Port. Multiple CDP devices can be neighbors on the same port if they are connected to the switch through a non-CDP device, such as some hubs.

CDP Version Data. The HP ProCurve switches use CDP-v1, but do not include IP prefix information, which is a router function; not a switch function.

Port Trunking with CDP. Where a static port trunk forms the link between the switch and another CDP device, only one physical link in the trunk is used to transmit outbound CDP packets.

CDP-Capable Hubs. Some hubs are capable of running CDP, but also forward CDP packets as if the hubs themselves were transparent to CDP. Such hubs will appear in the switch's CDP Neighbor table and will also maintain a CDP neighbor table similar to that for switches. For more information, refer to the documentation provided for the specific hub.

Clearing the Switch's CDP Neighbors Table. Use the switch's System Information screen to disable CDP on the switch and then re-enable it.

Troubleshooting CDP Operation

The switch does not appear in the CDP Neighbors table of an adjacent CDP Device. This may be due to any of the following:

- If there is more than one physical path between the switch and the other CDP device and STP is running on the switch, then STP will block the redundant link(s). In this case, the switch port on the remaining open link may not be a member of an untagged VLAN, or any untagged VLANs to which the port belongs may not have an IP address.
- The adjacent device's CDP Neighbors table may be full. Refer to the documentation provided for the adjacent CDP device to determine the table's capacity, and then view the device's Neighbors table to determine whether it is full.

One or more CDP neighbors appear intermittently or not at all in the switch's CDP Neighbors table. This may be caused by more than 60 neighboring devices sending CDP packets to the switch. Exceeding the 60-neighbor limit can occur, for example, where multiple neighbors are connected to the switch through non-CDP devices such as many hubs.

The Same CDP Switch or Router Appears on More Than One Port in the CDP Neighbors Table. Where CDP is running, a switch or router that is the STP root transmits outbound CDP packets over all links, including redundant links that STP may be blocking in non-root devices. In this case, the non-root device shows an entry in its CDP Neighbors table for every port on which it receives a CDP packet from the root device. See "Effect of Spanning Tree (STP) On CDP Packet Transmission" on page 37.

An IP Address of 127.0.0.1 Appears in the CDP Neighbors Table. This is the loopback IP address, which a port places in outbound CDP packets if none of the VLANs to which the port belongs has an IP address. (See "How the Switch Selects the IP Address To Include in Outbound CDP Packets" on page 38.)

New Time Synchronization Protocol Options

Using time synchronization ensures a uniform time among interoperating devices. This helps you to manage and troubleshoot switch operation by attaching meaningful time data to event and error messages.

Formerly, TimeP was the only time protocol available for time synchronization in the HP ProCurve Switches 1600M, 2400M, 2424M, 4000M, and 8000M. Beginning with software release C.09.xx, the switches also offer SNTP (Simple Network Time Protocol) and a new **Time Sync. mode** parameter for changing the time protocol selection (or turning off time protocol operation).

Notes

- Although you can create and save configurations for both time protocols without conflicts, the switch allows only one active time protocol at any time.
 - Time synchronization is no longer active in the factory default configuration. You must first select the desired protocol, and then enable it.
 - In the factory-default configuration for release C.09.xx and later, the time synchronization method is set to None. (In earlier releases, the default was TimeP with DHCP enabled for acquiring a TimeP server address).
 - If you configure SNTP operation in the switch, but later download a configuration created using a pre-C.09.xx version of the software, the SNTP configuration will be replaced by the non-SNTP time synchronization settings in the downloaded configuration file.
-

TimeP Time Synchronization

You can either manually assign the switch to use a TimeP server or use DHCP to assign the TimeP server. In either case, the switch can get its time synchronization updates from only one, designated Timep server.

SNTP Time Synchronization

SNTP provides two operating modes:

- **Broadcast Mode:** The switch acquires time updates by accepting the time value from the first SNTP time broadcast detected. (In this case, the SNTP server must be configured to broadcast time updates to the network broadcast address. Refer to the documentation
-

provided with your SNTP server application.) Once the switch detects a particular server, it ignores time broadcasts from other SNTP servers unless the configurable **Poll Interval** expires three consecutive times without an update received from the first-detected server.

Note

To use Broadcast mode, the switch and the SNTP server must be in the same broadcast domain.

- **Unicast Mode:** The switch requests a time update from the configured SNTP server. This option provides increased security over the Broadcast mode by specifying which time server to use instead of using the first one detected through a broadcast.

Overview: Selecting a Time Synchronization Protocol or Turning Off Time Protocol Operation

General Steps for Running a Time Protocol on the Switch:

1. Select the time synchronization protocol: **SNTP** or **TIMEP**.
2. Enable the protocol. The choices are:
 - SNTP: **Broadcast** or **Unicast**
 - TimeP: **DHCP** or **Manual**
3. Configure the remaining parameters for the time protocol you selected.

The switch retains the parameter settings for both time protocols even if you change from one protocol to the other. Thus, if you select a time protocol the switch uses the parameters you last configured for the selected protocol.

Note that simply selecting a time synchronization protocol does not enable that protocol on the switch unless you also enable the protocol itself (step 2, above).

Disabling Time Synchronization

To disable time synchronization without changing the Timep or SNTP configuration:

1. Go to the System Information screen of the Menu interface.
2. Set the **Time Sync. mode** parameter to **None**, then press [Enter], then [S] (for **Save**).

SNTP: Viewing, Selecting, and Configuring

Table 4. SNTP Parameters

SNTP Parameter	Operation
Time Sync Mode	Used to select either SNTP, TIMEP, or None as the time synchronization method.
SNTP Mode	
Disabled	The Default. SNTP does not operate, even if specified by the Menu interface Time Sync. mode parameter.
Unicast	Directs the switch to poll a specific server for SNTP time synchronization. Requires a server address.
Broadcast	Directs the switch to acquire its time synchronization from data broadcast by any SNTP server to the network broadcast address. The switch uses the first server detected and ignores any others. However, if the Poll Interval expires three times without the switch detecting a time update from the original server, the switch accepts a broadcast time update from the next server it detects.
Poll Interval (seconds)	In Unicast Mode: Specifies how often the switch polls the designated SNTP server for a time update. In Broadcast Mode: Specifies how often the switch polls the network broadcast address for a time update.
Server Address	Used only when the SNTP Mode is set to Unicast . Specifies the IP address of the SNTP server that the switch accesses for time synchronization updates.
Server Version	Default: 3; range: 1 - 7. Specifies the SNTP software version to use, and is assigned on a per-server basis. The version setting is backwards-compatible. For example, using version 3 means that the switch accepts versions 1 through 3.

Viewing and Configuring SNTP

To View, Enable, and Modify SNTP Time Protocol:

1. From the Main Menu, select:
 2. **Switch Management Access Configuration (IP, SNMP, Console) . . .**
 1. **IP Configuration**

```

                                DEFAULT_CONFIG
----- CONSOLE - MANAGER MODE -----
Switch Management Access Configuration - Internet (IP) Service

Time Sync. mode [None] : None ← Time Protocol Selection Parameter
                                - TIMEP
                                - SNTP
                                - None (Default)

IP Config [DHCP/Bootp] : DHCP/Bootp
IP Address :
Subnet Mask :
Gateway :

Actions->  Cancel   Edit   Save   Help

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

```

Figure 17. The Internet (IP) Service Screen (Default Values)

2. Press [E] (for **Edit**). The cursor moves to the **Time Sync. mode** field.
3. Use the Space bar to select **SNTP**, then press the downarrow key once to display and move to the **SNTP Mode** field.
4. Do one of the following:
 - Use the Space bar to select the **Broadcast** mode, then press the downarrow key to move the cursor to the **Poll Interval** field, and go to step 5. (For Broadcast mode details, see "SNTP Operating Modes" on page 42.)

```

Time Sync. mode [None] : SNTP
SNTP Mode [Disabled] : Broadcast
Poll Interval (sec) [720] : 720

```

- Use the Space bar to select the **Unicast** mode, then do the following:
 - i. Press the rightarrow key to move the cursor to the **Server Address** field.
 - ii. Enter the IP address of the SNTP server you want the switch to use for time synchronization.

Note: This step replaces any previously configured server IP address.

- iii. Press the downarrow key to move the cursor to the **Server Version** field. Enter the value that matches the SNTP server version running on the device you specified in the preceding step (step ii). If you are unsure which version to use, HP recommends leaving this value at the default setting of **3** and testing SNTP operation to determine whether any change is necessary.
- iv. Press the rightarrow key to move the cursor to the **Poll Interval** field, then go to step 5.

```

Time Sync. mode [None] : SNTP
SNTP Mode [Disabled] : Unicast           Server Address : 10.28.227.15
Poll Interval (sec) [720] : 720         Server Version [3] : 3

```

5. In the **Poll Interval** field, enter the time in seconds that you want for a Poll Interval. (For Poll Interval operation, see table 4, “SNTP Parameters”, on page 44.)
6. Press [Enter] to return to the Actions line, then [S] (for **Save**).

TimeP: Viewing, Selecting, and Configuring

Table 5. Timep Parameters

SNTP Parameter	Operation
Time Sync. mode	Used to select either TIMEP, SNTP, or None (the default) as the time synchronization method.
Timep Mode	
Disabled	The Default. Timep does not operate, even if specified by the Menu interface Time Sync. mode parameter.
DHCP	When Timep is selected as the time synchronization method, the switch attempts to acquire a Timep server IP address via DHCP. If the switch receives a server address, it polls the server for updates according to the Timep poll interval. If the switch does not receive a Timep server IP address, it cannot perform time synchronization updates.
Manual	When Timep is selected as the time synchronization method, the switch attempts to poll the specified server for updates according to the Timep poll interval. If the switch fails to receive updates from the server, time synchronization updates do not occur.
Server Address	Used only when the TimeP Mode is set to Manual . Specifies the IP address of the TimeP server that the switch accesses for time synchronization updates. You can configure one server.
Poll Interval (minutes)	Default: 720 minutes. Specifies the interval the switch waits between attempts to poll the TimeP server for updates.

Viewing and Configuring TimeP

To View, Enable, and Modify the TimeP Protocol:

1. From the Main Menu, select:
 2. **Switch Management Access Configuration (IP, SNMP, Console) ...**
 1. **IP Configuration**

```

                                DEFAULT_CONFIG
----- CONSOLE - MANAGER MODE -----
Switch Management Access Configuration - Internet (IP) Service

Time Sync. mode [None] : None ← Time Protocol Selection Parameter
                                - TIMEP
                                - SNTP
                                - None (the Default)

IP Config [DHCP/Bootp] : DHCP/Bootp
IP Address :
Subnet Mask :
Gateway :

Actions->  Cancel   Edit   Save   Help

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

Figure 18. The Internet (IP) Service Screen (Default Values)

2. Press [E] (for **Edit**). The cursor moves to the **Time Sync. mode** field.
3. Use the Space bar to select **TIMEP**, then press the downarrow key once to display and move to the **TimeP Mode** field.
4. Do one of the following:
 - Use the Space bar to select the **DHCP** mode, then press the downarrow key to move the cursor to the **Poll Interval** field, and go to step 5.

```

Time Sync. mode [None] : TIMEP
TimeP Mode [Disabled] : DHCP
Poll Interval (min) [720] : 720
```

- Use the Space bar to select the **Manual** mode.
 - i. Press the rightarrow key to move the cursor to the **Server Address** field.

- ii. Enter the IP address of the TimeP server you want the switch to use for time synchronization.

Note: This step replaces any previously configured TimeP server IP address.

- iii. Press the rightarrow key to move the cursor to the **Poll Interval** field, then go to step 5.

```
Time Sync Method [None] : TIMEP
TimeP Mode [Disabled] : Manual      Server Address : 10.28.227.141
Poll Interval (min) [720] : 720
Time Zone [0] : 0
Daylight Time Rule [None] : None
```

5. In the **Poll Interval** field, enter the time in minutes that you want for a TimeP Poll Interval.

Press [Enter] to return to the Actions line, then [S] (for **Save**) to enter the new time protocol configuration in both the startup-config and running-config files.

SNTP Messages in the Event Log

If an SNTP time change of more than three seconds occurs, the switch's event log records the change. SNTP time changes of less than three seconds do not appear in the Event Log.

Operation and Enhancements for Multimedia Traffic Control (IGMP)

How Data-Driven IGMP Operates

The information in this section supplements the information provided under "Multimedia Traffic Control with IP Multicast (IGMP)" beginning on page 9-91 in the Management and Configuration Guide included with your HP ProCurve Switch 1600M, 2400M, 2424M, 4000M, and 8000M, and also available at <http://www.hp.com/go/hpprocurve>.

This section uses the following terms to describe IGMP operation:

- **Querier:** A required IGMP device that facilitates the IGMP protocol and traffic flow on a given LAN. This device tracks which ports are connected to devices (IGMP clients) that belong to specific multicast groups, and triggers updates of this information. With IGMP enabled, the switch uses data from the queries to determine whether to forward or block multicast traffic on specific ports. When the switch has an IP address on a given VLAN, it automatically operates as a Querier for that VLAN if it does not detect a multicast router or another switch functioning as a Querier.
- **IGMP Device:** A switch or router running IGMP traffic control features.
- **IGMP Host:** An end-node device running an IGMP (multipoint, or multicast communication) application.

Without IGMP enabled, the switch simply floods all IP multicast traffic it receives on a given VLAN through all ports on that VLAN (except the port on which it received the traffic). This can result in significant and unnecessary bandwidth usage in networks where IP multicast traffic is a factor. Data-Driven IGMP reduces this problem by authorizing the switch to restrict multicast traffic only to ports where a given multicast group should flow.

An IP multicast packet includes the multicast group (address) to which the packet belongs. When an IGMP client connected to a switch port needs to receive multicast traffic from a specific group, it joins the group by sending an IGMP report (join request) to the network. (The multicast group specified in the join request is determined by the requesting application running on the IGMP client.) When a networking device with IGMP enabled receives the join request for a specific group, it forwards any IP multicast traffic it receives for that group through the port on which the join request was received. To reduce unnecessary traffic, the networking device does not forward a given group's

multicast packets to ports from which a join request for that group has not been received. (If the switch or router has not received any join requests for a given multicast group, it drops the traffic it receives for that group.)

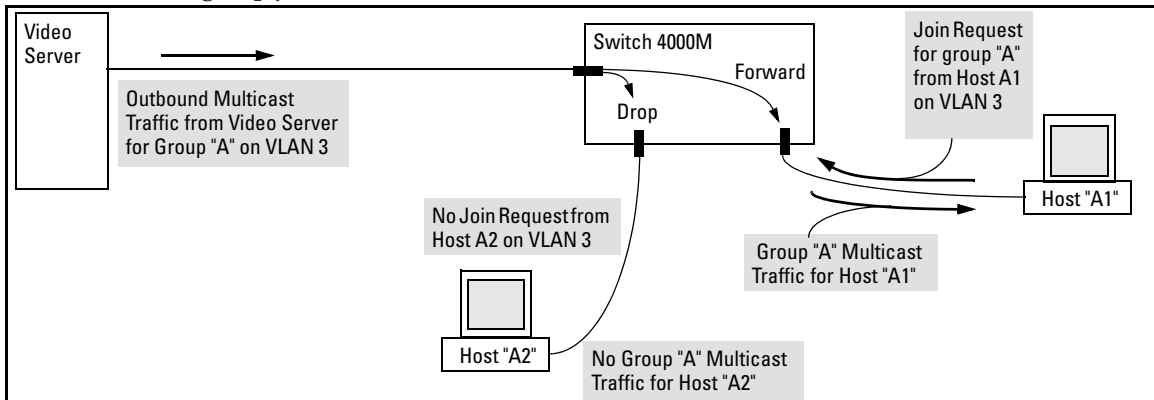


Figure 19. Example of Data-Driven IGMP Operation

Thus, after you enable IGMP on a VLAN configured in the switch, it continually listens for IGMP messages and IP multicast traffic on all ports in the VLAN, and forwards IGMP traffic for a given multicast address only through the port(s) on that VLAN where an IGMP report (join request) for that address was received from an IGMP client device.

Note

IP multicast traffic groups are identified by IP addresses in the range of 224.0.0.0 to 239.255.255.255.

Incoming IGMP packets intended for reserved, or "well-known" multicast addresses automatically flood through all ports (except the port on which the packets entered the switch). For more on this topic, see "The Switch Excludes Well-Known or Reserved Multicast Addresses from IP Multicast Filtering" on page 59.

New: IGMP Now Operates With or Without IP Addressing

Formerly, IGMP operation required an IP address and subnet mask for each VLAN running IGMP. Beginning with release C.09.xx, you can configure IGMP on VLANs that do not have IP addressing. The benefit of IGMP without IP addressing is a reduction in the number of IP addresses you have to use and configure. This can be significant in a network with a large number of VLANs. The limitation on IGMP without IP addressing is that the switch cannot become Querier on any VLANs for which it has no IP address—so the network administrator must ensure that another IGMP device will act as Querier. It is also advisable to have an additional IGMP device available as a backup Querier.

IGMP Function Available With IP Addressing Configured on the VLAN	Available Without IP Addressing?	Operating Differences Without an IP Address
Drop multicast group traffic for which there have been no join requests from IGMP clients connected to ports on the VLAN.	Yes	None
Forward multicast group traffic to any port on the VLAN that has received a join request for that multicast group.	Yes	None
Forward join requests (reports) to the Querier.	Yes	None
Configure individual ports in the VLAN to Auto (the default)/ Blocked , or Forward .	Yes	None
Configure IGMP traffic forwarding to normal or high-priority forwarding.	Yes	None
Age-Out IGMP group addresses when the last IGMP client on a port in the VLAN leaves the group.	Yes	Requires that another IGMP device in the VLAN has an IP address and can operate as Querier. This can be a multi-cast router or another switch configured for IGMP operation. (HP recommends that the VLAN also include a device operating as a backup Querier in case the device operating as the primary Querier fails for any reason.
Support Fast-Leave IGMP (below) and Forced Fast-Leave IGMP (page 53).	Yes	
Support automatic Querier election.	No	Querier operation not available.
Operate as the Querier.	No	Querier operation not available.
Available as a backup Querier.	No	Querier operation not available.

Fast-Leave IGMP

IGMP Operation Presents a "Delayed Leave" Problem. Where multiple IGMP clients are connected to the same port on an IGMP device (switch or router), if only one IGMP client joins a given multicast group, then later sends a Leave Group message and ceases to belong to that group, the IGMP device retains that IGMP client in its IGMP table and continues forwarding IGMP traffic to the IGMP client until the Querier triggers confirmation that no other group members exist on the same port. This means that the switch continues to transmit unnecessary multicast traffic through the port until the Querier renews multicast group status.

Fast-Leave IGMP Reduces Leave Delays. Fast-Leave IGMP automatically operates on a port if an IGMP client connects to the port and there are no other end nodes detected on that port. In this case, when the client leaves a multicast group, Fast-Leave IGMP automatically accelerates the blocking of further, unnecessary multicast traffic from that group to the former IGMP client. This improves performance by reducing the amount of multicast traffic going through the port to the IGMP client after the client leaves a multicast group.

Automatic Fast-Leave Operation. If a switch port is :

- a. Connected to only one end node
- b. The end node currently belongs to a multicast group; i.e. is an IGMP client
- c. The end node subsequently leaves the multicast group

Then the switch does not need to wait for the Querier status update interval, but instead immediately removes the IGMP client from its IGMP table and ceases transmitting IGMP traffic to the client. (If the switch detects multiple end nodes on the port, automatic Fast-Leave does not activate—regardless of whether one or more of these end nodes are IGMP clients.)

In the next figure, automatic Fast-Leave operates on the switch ports for IGMP clients "3A" and "5B", but not on the switch port for IGMP clients "7A" and 7B, Server "7C", and printer "7D".




Figure 20. Example of Automatic Fast-Leave IGMP Criteria

When client "3A" running IGMP is ready to leave the multicast group, it transmits a Leave Group message. Because the switch knows that there is only one end node on port 3, it removes the client from its IGMP table and halts multicast traffic (for that group) to port 3. If the switch is not the Querier, it does not wait for the actual Querier to verify that there are no other group members on port 3. If the switch itself is the Querier, it does not query port 3 for the presence of other group members.

Note that Fast-Leave operation does not distinguish between end nodes on the same port that belong to different VLANs. Thus, for example, even if all of the devices on port 7 in figure 20 belong to different VLANs, Fast-Leave does not operate on port 7.

New: Forced Fast-Leave IGMP

Forced Fast-Leave IGMP speeds up the process of blocking unnecessary IGMP traffic to a switch port that is connected to multiple end nodes. (This feature does not activate on ports where the switch detects only one end node. Instead, the regular Fast Leave described in the preceding section activates.) For example, in figure 20, even if you configured Forced Fast-Leave on all ports in the switch, the feature would activate only on port 7 (which has multiple end nodes) when a Leave Group request arrived on that port.

When a port having multiple end nodes receives a Leave Group request from one end node for a given multicast group "X", Forced Fast-Leave activates and waits a small amount of time to receive a join request from any other group "X" member on that port. If the port does not receive a join request for that group within the forced-leave interval, the switch then blocks any further group "X" traffic to the port.

Configuration Options for Forced Fast-Leave

Feature	Default	Settings	Function
Forced Fast-Leave state	2 (disabled)	1 (enabled) 2 (disabled)	Uses the setmib command to enable or disable Forced Fast-Leave on individual ports. When enabled on a port, Forced Fast-Leave operates only if the switch detects multiple end nodes (and at least one IGMP client) on that port.

Note on VLAN Numbers:

In the HP ProCurve Switches 1600M, 2400M, 2424M, 4000M, and 8000M, the **walkmib** and **setmib** commands use an internal VLAN number (and not the VLAN ID, or VID) to display or change many per-vlan features, such as the Forced Fast-Leave state. Because the internal VLAN number for the default VLAN is always 1 (regardless of whether VLANs are enabled on the switch), and because a discussion of internal VLAN numbers for multiple VLANs is beyond the scope of this document, the discussion here concentrates on examples that use the default VLAN.

Listing the Forced Fast-Leave Configuration

The Forced Fast-Leave configuration data is available in the switch's MIB (Management Information Base), and includes the state (enabled or disabled) for each port and the Forced-Leave Interval for all ports on the switch.

To List the Forced Fast-Leave State for all Ports in the Switch. Go to the switch's command prompt and use the **walkmib** command, as shown below.

1. From the Main Menu, select:

5. Diagnostics . . .

4. Command Prompt

2. Do one of the following:

- If VLANs are not enabled on the switch, go to step 3.
- If VLANs are enabled on the switch:
 - i. You will be prompted to select a VLAN. For example:

```
Select VLAN : DEFAULT VLAN
```

- ii. Because you can list the Forced Fast-Leave state for all ports on the switch from any VLAN, just press [Enter] to select the displayed VLAN.

3. Enter either of the following walkmib command options:

```
walkmib hpSwitchIcmpPortForcedLeaveState
```

- OR -

```
walkmib 1.3.6.1.4.1.11.2.14.11.5.1.7.1.15.3.1.5
```

The resulting display lists the Forced Fast-Leave state for all ports in the switch, by VLAN. (A port belonging to more than one VLAN will be listed once for each VLAN, and if multiple VLANs are *not* configured, all ports will be listed as members of the default VLAN.) The following command produces a listing such as that shown in figure 21:

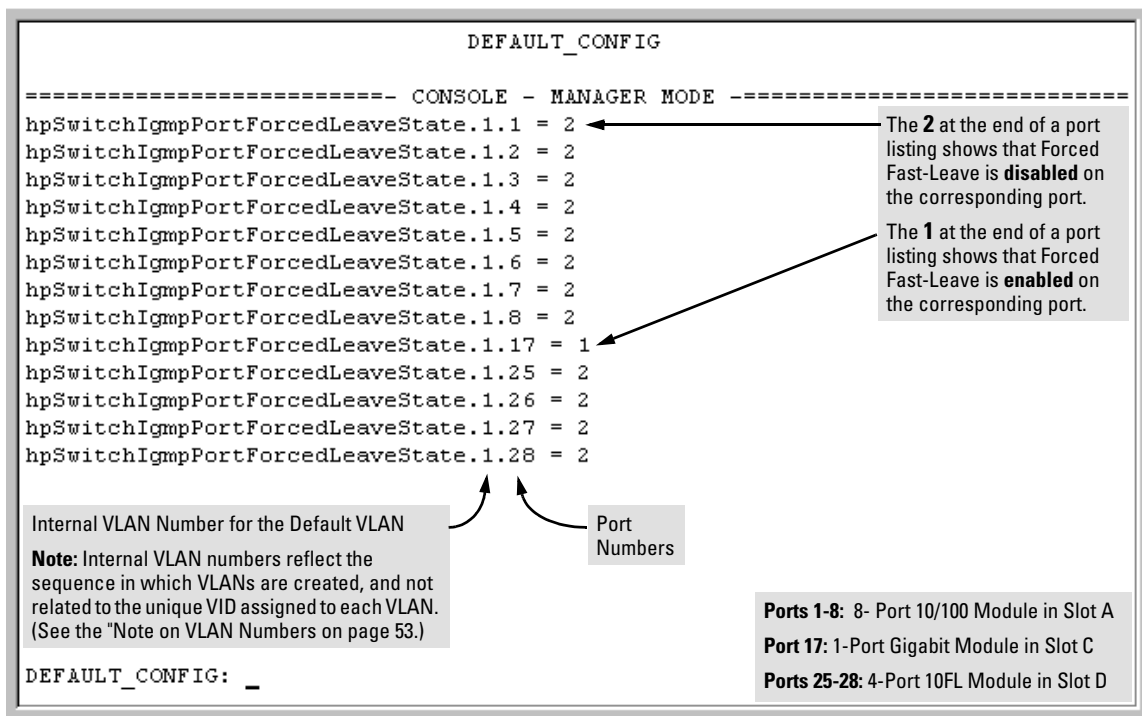


Figure 21. Example of a Forced Fast-Leave Listing where all Ports are Members of the Default VLAN

To List the Forced Fast-Leave State for a Single Port. (See the "Note on VLAN Numbers" on page 53.)

Go to the switch's command prompt and use the **getmib** command, as shown below.

- From the Main Menu, select:

5. Diagnostics . . .

4. Command Prompt

- Enter either of the following walkmib command options:

```
getmib hpSwitchIcmpPortForcedLeaveState.<vlan number><.port number>
```

- OR -

```
getmib 1.3.6.1.4.1.11.2.14.11.5.1.7.1.15.3.1.5.<vlan number><.port number>
```

For example, the following command to list the state for port 7 (which, in this case, belongs to the default VLAN) produces a listing similar to that shown in figure 22:

DEFAULT_CONFIG: setmib hpSwitchIcmpPortForcedLeaveState.1.17 -i 1

```
DEFAULT_CONFIG
----- CONSOLE - MANAGER MODE -----
hpSwitchIcmpPortForcedLeaveState.1.17 = 1
DEFAULT_CONFIG: _
```

Verifies Forced Fast-Leave enabled.
17 indicates port 17.
1 indicates the default VLAN. (See the note on page 53.)

Figure 23. Example of Changing the Forced Fast-Leave Configuration on Port 17

Using the Switch as Querier

Querier Operation

The function of the IGMP Querier is to poll other IGMP-enabled devices in an IGMP-enabled VLAN to elicit group membership information. The switch performs this function if there is no other device in the VLAN, such as a multicast router, to act as Querier. Although the switch automatically ceases Querier operation in an IGMP-enabled VLAN if it detects another Querier on the VLAN, you can also use the Command Prompt to disable the Querier function for that VLAN.

Note

A Querier is required for proper IGMP operation. For this reason, if you disable the Querier function on a the switch, ensure that there is an IGMP Querier (and, preferably, a backup Querier) available on the same VLAN.

If the switch becomes the Querier for a particular VLAN (for example, the DEFAULT_VLAN), then subsequently detects queries transmitted from another device on the same VLAN, the switch ceases to operate as the Querier for that VLAN. If this occurs, the switch Event Log lists a pair of messages similar to these:

```
I 01/15/01 09:01:13 igmp: DEFAULT_VLAN: Other Querier detected
I 01/15/01 09:01:13 igmp: DEFAULT_VLAN: This switch is no longer Querier
```

In the above scenario, if the other device ceases to operate as a Querier on the default VLAN, then the switch detects this change and can become the Querier as long as it is not pre-empted by some other IGMP Querier on the VLAN. In this case, the switch Event Log lists messages similar to the following to indicate that the switch has become the Querier on the VLAN:

I 01/15/01 09:21:55 igmp: DEFAULT_VLAN: Querier Election in process
I 01/15/01 09:22:00 igmp: DEFAULT_VLAN: This switch has been elected as Querier

Changing the Querier Configuration Setting

The Querier feature, by default, is enabled and in most cases should be left in this setting. If you need to change the querier setting, you can do so using the IGMP Configuration MIB. To disable the querier setting, select the Command Prompt from the Diagnostics Menu and enter this command:

```
setmib hpSwitchIgmpQuerierState.<vlan number> -i 2
```

To enable the querier setting, select the Command Prompt from the Main Menu and enter this command:

```
setmib hpSwitchIgmpQuerierState.<vlan number> -i 1
```

To view the current querier setting, select the Command Prompt from the Main Menu and enter this command:

```
getmib hpSwitchIgmpQuerierState.<vlan number>
```

where:

<vlan number> is the sequential (index) number of the specific VLAN. If no VLANs are configured, use "1". For example:

```
getmib hpSwitchIgmpQuerierState.1
```

The Switch Excludes Well-Known or Reserved Multicast Addresses from IP Multicast Filtering

Each multicast host group is identified by a single IP address in the range of 224.0.0.0 through 239.255.255.255. Specific groups of consecutive addresses in this range are termed "well-known" addresses and are reserved for predefined host groups. IGMP does not filter these addresses, so any packets the switch receives for such addresses are flooded out all ports assigned to the VLAN on which they were received (except the port on which the packets entered the VLAN). The following table lists the 32 well-known address groups (8192 total addresses) that IGMP does not filter on the 1600M, 2400M, 2424M, 4000M, and 8000M switches.

Table 6. Well-Known IP Multicast Address Groups Excluded from IGMP Filtering

Groups of Consecutive Addresses in the Range of 224.0.0.X to 239.0.0.X*		Groups of Consecutive Addresses in the Range of 224.128.0.X to 239.128.0.X*	
224.0.0.x	232.0.0.x	224.128.0.x	232.128.0.x
225.0.0.x	233.0.0.x	225.128.0.x	233.128.0.x
226.0.0.x	234.0.0.x	226.128.0.x	234.128.0.x
227.0.0.x	235.0.0.x	227.128.0.x	235.128.0.x
228.0.0.x	236.0.0.x	228.128.0.x	236.128.0.x
229.0.0.x	237.0.0.x	229.128.0.x	237.128.0.x
230.0.0.x	238.0.0.x	230.128.0.x	238.128.0.x
231.0.0.x	239.0.0.x	231.128.0.x	239.128.0.x

* X is any value from 0 to 255.

New: Menu Enhancement for Moving from Operator Access to Manager Access

Prior to release C.09.*xxx*, with both the Operator (read-only) and Manager (read-write) password configured, if you entered the menu interface at the Operator level you had to log out and re-enter the menu interface to move to the Manager level. Now, using the new **Enable Manager Mode** option in the Operator-level Main menu, you can move directly to the Manager level by entering the correct Manager password. To do so, go to the Main Menu, select list item 4, Enable Manager Mode, and, when prompted, enter the appropriate password. If TACACS+ authentication is operating for read-write access, you will be prompted for the Enable-Level password configured for the switch in the TACACS+ server.

```

                                DEFAULT_CONFIG
----- CONSOLE - OPERATOR MODE -----
                                Main Menu

1. Status and Counters...
2. Event Log
3. Diagnostics...
4. Enable Manager Mode...
0. Logout

Displays menu of diagnostics tools, including the command prompt.
To select menu item, press item number, or highlight item and press <Enter>.
```


The image shows a terminal window displaying a menu. The menu title is 'DEFAULT_CONFIG' followed by a separator line '----- CONSOLE - OPERATOR MODE -----' and the subtitle 'Main Menu'. The menu items are: 1. Status and Counters..., 2. Event Log, 3. Diagnostics..., 4. Enable Manager Mode..., and 0. Logout. The item '4. Enable Manager Mode...' is highlighted with a black background. A grey callout box with an arrow points to this item, containing the text: 'New Enable Manager option for moving from Operator (read-only) to Manager (read-write) access.' At the bottom of the terminal window, there is a black bar with white text: 'Displays menu of diagnostics tools, including the command prompt. To select menu item, press item number, or highlight item and press <Enter>.'

Figure 24. The Main Menu with the "Enable Manager Mode" Option

Configuring and Using HP ProCurve Stack Management

HP ProCurve Stack Management (termed *stacking* in this document) enables you to use a single IP address and standard network cabling to manage a group of up to 16 switches in the same subnet (broadcast domain). Using stacking, you can:

- Reduce the number of IP addresses needed in your network.
- Simplify management of small workgroups or wiring closets while scaling your network to handle increased bandwidth demand.
- Eliminate any specialized cables for stacking connectivity and remove the distance barriers that typically limit your topology options when using other stacking technologies.
- Add switches to your network without having to first perform IP addressing tasks.

Components of HP ProCurve Stack Management

Table 7. Stacking Definitions

Stack	Consists of a Commander switch and any Member switches belonging to that Commander.
Commander	A switch that has been manually configured as the controlling device for a stack. When this occurs, the switch's stacking configuration appears as Commander .
Candidate	A switch that is ready to join (become a Member of) a stack through either automatic or manual methods. A switch configured as a Candidate is not in a stack.
Member	The switch that has joined a stack and is under the control of the stack's Commander.




Figure 25. Illustration of a Switch Moving from Candidate to Member

General Operation

After you configure one switch to operate as the Commander of a stack, additional switches can join the stack by either automatic or manual methods. After a switch becomes a Member, you can work through the Commander switch to further configure the Member switch as necessary for all of the additional software features available in the switch.

The Commander switch serves as the entry point for access to the Member switches. For example, the Commander's IP address becomes the path to a Member. The Commander's Manager password controls access to a Member.

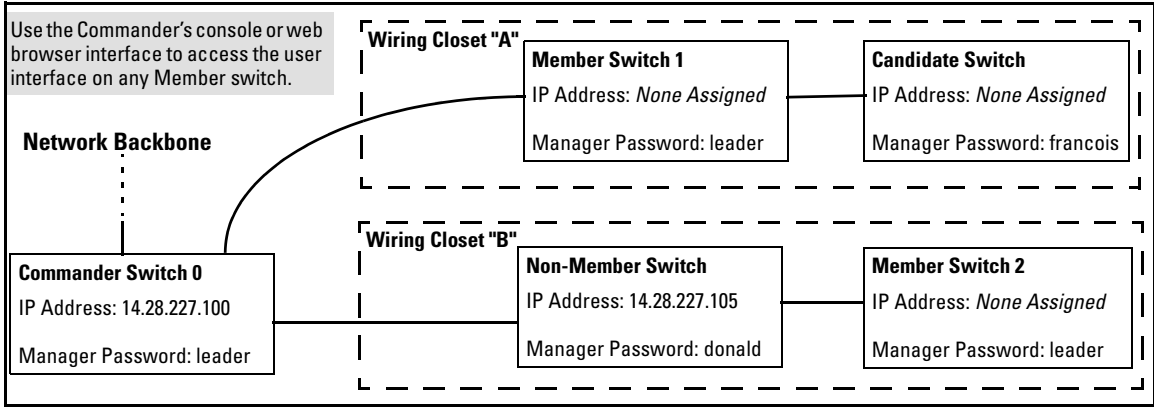


Figure 26. Example of Stacking with One Commander Controlling Access to Wiring Closet Switches

Interface Options. You can configure stacking through either the console menu interface or the web browser interface. For information on how to use the web browser interface to configure stacking, see the online Help for the web browser interface.

Changes to the Web Browser Interface for Commander Switches. Updating an HP ProCurve switch to software release C.08.xx or later and configuring the switch as a Commander for stacking introduces a modified web browser interface that differs in appearance from the version documented in the *Management and Configuration Guide* you received with your switches. Note that this change does not appear in the web browser interface for Candidate and Member switches, and switches on which the Stacking option is disabled. See page 84.

Changes to the Console Interface. Updating an HP ProCurve switch to software release C.08.xx or later adds the Stacking option as item 8 in the Main Menu.

```

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

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

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

```

Figure 27. Main Menu in Software Release C.08.xx (or Later)

Operating Rules for Stacking

General Rules

- Stacking is an optional feature (enabled in the default configuration) and can easily be disabled. Stacking has no effect on the normal operation of the switch in your network.
- A stack requires one Commander switch. (Only one Commander allowed per stack.)
- All switches in a particular stack must be in the same subnet (broadcast domain). A stack cannot cross a router.
- A stack accepts up to 16 switches (numbered 0-15), including the Commander (always numbered 0).
- There is no limit on the number of stacks in the same subnet (broadcast domain), however a switch can belong to only one stack.
- If VLANs are enabled, stacking uses only the default VLAN on any switch, even if you change the name and/or ID number of the default VLAN. (See "Stacking Operation with a Tagged VLAN" on page 83.)
- Stacking allows intermediate devices that do not support stacking. This enables you to include devices that are distant from the Commander.




Figure 28. Example of a Non-Stacking Device Used in a Stacking Environment

Specific Rules for Commander, Candidate, and Member Switches

	IP Addressing and Stack Name	Number Allowed Per Stack	Passwords	SNMP Communities
Commander	<p>IP Addr: Requires an assigned IP address and mask for access via the network. (Otherwise an IP address is optional.)</p> <p>Stack Name: Required</p>	Only one Commander switch is allowed per stack.	<p>The Commander's Manager and Operator passwords are assigned to any switch becoming a Member of the stack.</p> <p>If you change the Commander's passwords, the Commander propagates the new passwords to all stack Members.</p>	Standard SNMP community operation. The Commander also operates as an SNMP proxy to Members for all SNMP communities configured in the Commander.
Candidate	<p>IP Addr: Optional. Configuring an IP address allows access via Telnet or web browser interface while the switch is not a stack member. In the factory default configuration the switch automatically acquires an IP address if your network includes DHCP service.</p> <p>Stack Name: N/A</p>	No limit.	<p>Passwords optional. If the Candidate becomes a stack Member, it assumes the Commander's Manager and Operator passwords.</p>	Uses standard SNMP community operation if the Candidate has its own IP addressing.
Member	<p>IP Addr: Optional. Configuring an IP address allows access via Telnet or web browser interface without going through the Commander switch. This is useful, for example, if the stack Commander fails and you need to convert a Member switch to operate as a replacement Commander.</p> <p>Stack Name: N/A</p>	Up to 15 Members per stack.	<p>When the switch joins the stack, it automatically assumes the Commander's Manager and Operator passwords and discards any passwords it may have had while a Candidate.</p> <p>Note: If a Member leaves a stack for any reason, it retains the passwords assigned to the stack Commander at the time of departure from the stack.</p>	Belongs to the same SNMP communities as the Commander (which serves as an SNMP proxy to the Member for communities to which the Commander belongs). To join other communities that <i>exclude</i> the Commander, the Member must have its own IP addressing. Loss of stack membership means loss of membership in any community that is configured only in the Commander. See "SNMP Community Operation in a stack" on 82.

Note

In the default stack configuration, the Candidate **Auto Join** parameter is enabled, but the Commander **Auto Grab** parameter is disabled. This prevents Candidates from automatically joining a stack prematurely or joining the wrong stack (if more than one stack Commander is configured in a subnet or broadcast domain). If you plan to install more than one stack in a subnet, HP recommends that you leave **Auto Grab** disabled on all Commander switches and manually add Members to their stacks. Similarly, if you plan to install a stack in a subnet (broadcast domain) where software release C.08.xx or later is running on switches not intended for stack membership, you should set the **Stack State** parameter (in the Stack Configuration screen) to **Disabled** on those particular switches.

Configuring and Bringing Up a Stack

This process assumes that:

- You have downloaded software version C.08.xx (or later) to all switches that you want to include in a stack. (You can get a copy of the software from HP's ProCurve web site and/or copy it from one switch to another. For downloading instructions, see appendix A, "File Transfers", in the *Management and Configuration Guide* you received with your switches.)
- All switches you want to include in a stack are connected to the same subnet.
- If VLANs are enabled on the switches comprising the stack, then the ports linking the stacked switches must be on the default VLAN in each switch. If the default VLAN is tagged, then it must have the same VLAN ID (VID) for each port providing a link to the stack. (See "Stacking Operation with a Tagged VLAN" on page 83.)

Depending on how Commander and Candidate switches are configured, Candidates can join a stack either automatically or by manual assignment through the Commander. The following table shows your control options for adding Members to a stack.

Table 8. Stacking Configuration Guide

Join Method	Commander (IP Addressing Required)	Candidate (IP Addressing Optional)	
	Auto Grab	Auto Join	Passwords
Automatically add Candidate to stack	Yes	Yes (default)	No (default)*
Manually add Candidate to stack: (Prevent automatic joining of switches you don't want in the stack)	No (default)	Yes (default)	Optional*
	Yes	No	Optional*
	Yes	Yes (default) or No	Configured
Prevent a switch from being a Candidate	N/A	Disabled	Optional

*The Commander's Manager and Operator passwords propagate to all Members when they individually join the stack.

The easiest way to *automatically* create a stack is to:

1. Configure a switch as a Commander.
2. Configure IP addressing and a stack name on the Commander.
3. Set the Commander's Auto Grab parameter to Yes.
4. Connect Candidate switches (in their factory default configuration) to the network.

This approach automatically creates a stack of up to 16 switches (including the Commander). However this replaces manual control with an automatic process that may bring switches into the stack that you did not intend to include. With the Commander's **Auto Grab** parameter set to **Yes**, *any switch* conforming to all three of the following factors automatically becomes a stack Member:

- Default stacking configuration (**Stack State** set to **Candidate**, and **Auto Join** set to **Yes**)
- Same subnet (broadcast domain) and default VLAN as the Commander (If VLANs are used in the stack environment, see "Stacking Operation with a Tagged VLAN" on page 83.)
- No Manager password

Overview of How To Create a Stack

This section describes the general stack creation process. For the detailed configuration processes, see pages 68 through 77.

1. Determine the naming conventions for the stack. You will need a stack name. Also, to help distinguish one switch from another in the stack, you can configure a unique system name for each switch. Otherwise, the system name for a switch appearing in the Stacking Status screen appears as the stack name plus an automatically assigned switch number. For example:

```
Pacific Ocean
-----
CONSOLE - MANAGER MODE -----
Stacking - Stacking Status (All)

Stack Name      MAC Address      System Name      Status
-----
Big Waters      0060b0-880a80    Pacific Ocean    Commander Up
                0060b0-df1a00    Coral Sea        Member Up
Online          0060b0-df7680    online-0         Commander Up
                001083-3c7480    online-1         Member Up
                0060b0-312f00    online-2         Member Up
                001083-3c09c0    online-3         Member Up

Actions->  Back  Next page  Prev page  Help
Return to previous screen.
Use up/down arrow keys to scroll to other entries, left/right arrow keys to
change action selection, and <Enter> to execute action.
```

For status descriptions, see the table on page 83.

Stack with unique system name for each switch.

Stack named "Online" with no previously configured system names assigned to individual switches.

Figure 29. Use of System Name to Help Identify Individual Switches

2. Configure the Commander switch. Doing this first helps to establish consistency in your stack configuration, which can help prevent startup problems.
 - The Commander assigns its Manager and Operator passwords to any Candidate switch that joins the stack.
 - SNMP community names used in the Commander apply to stack members.
3. If you need to access Candidate switches through your network before they join the stack, assign IP addresses to these devices. Otherwise, IP addressing is optional for Candidates and Members.
4. Make a record of any Manager passwords assigned to the switches (intended for your stack) that are not currently members. (You will have to use these passwords to enable the protected switches to join the stack.)
5. If you are using VLANs in the stacking environment, you must use the default VLAN for stacking links. For more information, see "Stacking Operation with a Tagged VLAN" on page 83.
6. Ensure that all switches intended for the stack are connected to the same subnet (broadcast domain). As soon as you connect the Commander, it will begin discovering the available Candidates in the subnet.
 - If you configured the Commander to automatically add Members (**Auto Grab** set to **Yes**), then any discovered Candidates meeting the following criteria will automatically become stack Members:
 - **Auto Join** parameter set to **Yes** (the default)
 - Manager password not configured
 - If you configured the Commander to manually add Members (**Auto Grab** set to **No**—the default), you can begin the process of selecting and adding the desired Candidates.
7. Ensure that all switches intended for the stack have joined.
8. If you need to perform specific configuration or monitoring tasks on a Member, use the console interface on the Commander to select and access the Member.

Configuring a Commander Switch

A stack requires one Commander switch. If you plan to implement more than one stack in a subnet (broadcast domain), the easiest way to avoid unintentionally adding a Candidate to the wrong stack is to manually control the joining process by leaving the Commander's **Auto Grab** parameter set to **No** (the default).

1. Configure an IP address and subnet mask on the Commander switch.
2. Display the Stacking Menu by selecting **Stacking** in the Main Menu.

```

                                DEFAULT_CONFIG
----- CONSOLE - MANAGER MODE -----
                                Stacking Menu

1. Stacking Status (This Switch)
2. Stacking Status (All)
3. Stack Configuration
0. Return to Main Menu...

Shows the status of Stack.
To select menu item, press item number, or highlight item and press <Enter>.

```

Figure 30. The Default Stacking Menu

3. Display the Stack Configuration menu by pressing [3] to select **Stack Configuration**.

```

                                DEFAULT_CONFIG
----- CONSOLE - MANAGER MODE -----
                                Stacking - Stack Configuration

Stack State : Candidate
Auto Join [Yes] : Yes
Transmission Interval [60] : 60

Actions->  Cancel  Edit  Save  Help

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

```

Figure 31. The Default Stack Configuration Screen

4. Move the cursor to the Stack State field by pressing [E] (for **Edit**). Then use the Space bar to select the **Commander** option.
5. Press the downarrow key to display the Commander configuration fields in the Stack Configuration screen.

```

                                DEFAULT_CONFIG
=====-- CONSOLE - MANAGER MODE -=====
                                Stacking - Stack Configuration

Stack State : Commander
Stack Name : ██████████
Auto Grab [No] : No
Transmission Interval [60] : 60

Actions->  Cancel      Edit      Save      Help

████████████████████████████████████████████████████████████████████████████████
Use arrow keys to change field selection, <Space> to toggle field choices,
and <Enter> to go to Actions.

```

Figure 32. The Default Commander Configuration in the Stack Configuration Screen

6. Enter a unique stack name (up to 15 characters; no spaces) and press the downarrow key.
7. Ensure that you have a correct the **Auto Grab** setting, then press the downarrow key:
 - **No** (the default) prevents automatic joining by Candidates having their **Auto Join** set to **Yes**.
 - **Yes** enables the Commander to automatically take a Candidate into the stack as a Member if the Candidate has **Auto Join** set to **Yes** (the default Candidate setting).
8. Accept or change the transmission interval (default: 60 seconds), then press [Enter] to return the cursor to the **Actions** line.
9. Press [S] (for **S**ave) to save your configuration changes and return to the Stacking menu.

Your Commander switch should now be ready to automatically or manually acquire Member switches from the Candidate list, depending on your configuration choices.

Modifying or Disabling Stacking On a Candidate Switch

In its default stacking configuration, a Candidate switch can either automatically join a stack or be manually added ("pulled") into a stack by a Commander, depending on the Commander's **Auto Grab** setting. You can also reconfigure a Candidate switch to either "push" the Candidate into membership with a specific Commander's stack, convert the Candidate to a stack Commander (for a stack that does not already have a Commander), or to operate as a standalone switch without stacking. You can also change a Candidate's **Auto Join** or **Transmission Interval** settings. The following table lists the options:

Table 9. Candidate Configuration Options

Parameter	Default Setting	Other Settings
Stack State	Candidate	Commander, Member, or Disabled
Auto Join	Yes	No
Transmission Interval	60 Seconds	Range: 1 to 300 seconds

Use Telnet or the web browser interface to access the Candidate if it has an IP address. Otherwise, use a direct connection from a terminal device to the switch's console port. (For information on how to use the web browser interface, see the online Help provided for the browser.)

1. Display the Stacking Menu by selecting **Stacking** in the console Main Menu.
2. Display the Stack Configuration menu by pressing [3] to select **Stack Configuration**.




Figure 33. The Default Stack Configuration Screen

3. Move the cursor to the Stack State field by pressing [E] (for **Edit**).
4. Do one of the following:
 - To disable stacking on the Candidate, use the Space bar to select the **Disabled** option, then go to step 5.
Note: Disabling stacking on a Candidate removes the Candidate from all stacking menus.
 - To insert the Candidate into a specific Commander's stack:
 - i. Use the space bar to select Member.
 - ii. Press [Tab] once to display the **Commander MAC Address** parameter, then enter the MAC address of the desired Commander.
 - To change **Auto Join** or **Transmission Interval**, use [Tab] to select the desired parameter, and:
 - To change **Auto Join**, use the Space bar.
 - To change **Transmission Interval**, type in the new value in the range of 1 to 300 seconds.

Then go to step 5.

5. press [Enter] to return the cursor to the **Actions** line.

6. Press [S] (for **Save**) to save your configuration changes and return to the Stacking menu.

Manually Adding a Candidate to a Stack

In the default configuration, you must manually add stack Members from the Candidate pool. Reasons for a switch remaining a Candidate instead of becoming a Member include any of the following:

- **Auto Grab** in the Commander is set to **No** (the default).
- **Auto Join** in the Candidate is set to **No**.

Note: When a switch leaves a stack and returns to Candidate status, its **Auto Join** parameter resets to **No** so that it will not immediately rejoin a stack from which it has just departed.

- A Manager password is set in the Candidate.
- The stack is full.

Unless the stack is already full, you can use the Stack Management screen to manually convert a Candidate to a Member. If the Candidate has a Manager password, you will need to use it to make the Candidate a Member of the stack.

1. To add a Member, start at the Main Menu and select:

8. Stacking...

4. Stack Management

You will then see the Stack Management screen:

```
Pacific Ocean
----- CONSOLE - MANAGER MODE -----
Stacking - Stack Management

SN      MAC Address      System Name      Device Type      Status
-----
1      0060b0-df1a00      Coral Sea        HP 8000M         Member Up
2      080009-8c5080      North Atlantic   HP 8000M         Member Up

Actions->  Back  Add  Edit  Delete  Help

Return to previous screen.
Use up/down arrow keys to change record selection, left/right arrow keys to
change action selection, and <Enter> to execute action.
```

For status descriptions, see the table on page 83.

Figure 34. Example of the Stack Management Screen

2. Press [A] (for **Add**) to add a Candidate. You will then see this screen listing the available Candidates:

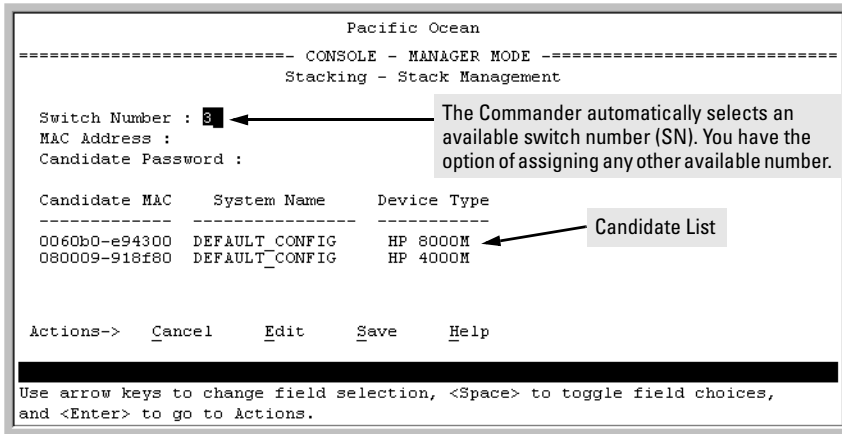


Figure 35. Example of Candidate List in Stack Management Screen

3. Either accept the displayed switch number or enter another available number. (The range is 0 - 15, with 0 reserved for the Commander.)
4. Use the downarrow key to move the cursor to the MAC Address field, then type the MAC address of the desired Candidate from the Candidate list in the lower part of the screen.
5. Do one of the following:
 - If the desired Candidate has a Manager password, press the downarrow key to move the cursor to the Candidate Password field, then type the password.
 - If the desired Candidate does not have a password, go to step 6.
6. Press [Enter] to return to the Actions line, then press [S] (for **S**ave) to complete the Add process for the selected Candidate. You will then see a screen similar to the one in figure 34 on page 72, with the newly added Member listed.

Note: If the message **Unable to add stack member: Invalid Password** appears in the console menu's Help line, then you either omitted the Candidate's Manager password or incorrectly entered the Manager password.

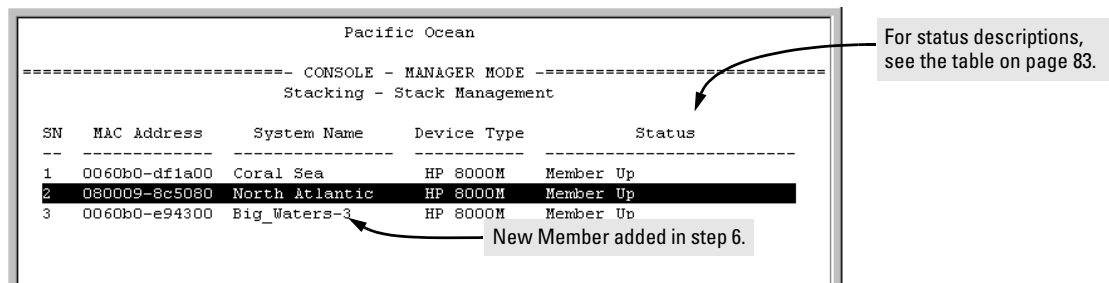


Figure 36. Example of Stack Management Screen After New Member Added

Moving a Member From One Stack to Another

Where two or more stacks exist in the same subnet (broadcast domain), you can easily move a Member of one stack to another stack if the destination stack is not full. (If you are using VLANs in your stack environment, see "Stacking Operation with a Tagged VLAN" on page 83.) This procedure is nearly identical to manually adding a Candidate to a stack (page 72). (If the stack from which you want to move the Member has a Manager password, you will need to know the password to make the move.)

1. To move a Member from one stack to another, go to the Main Menu of the Commander in the destination stack and display the Stacking Menu by selecting

8. Stacking...

2. To learn or verify the MAC address of the Member you want to move, display a listing of all Commanders, Members, and Candidates in the subnet by selecting:

2. Stacking Status (All)

You will then see the Stacking Status (All) screen:

```
Pacific Ocean
----- CONSOLE - MANAGER MODE -----
Stacking - Stacking Status (All)

Stack Name      MAC Address      System Name      Status
-----
Big Waters      0060b0-880a80   Pacific Ocean   Commander Up
Newstack        0060b0-df1a00   Coral Sea       Member Up
                080009-8c5080   North Atlantic  Member Up
                001083-c3fc00   Newstack-0      Commander Up
                080009-918f80   Newstack-1      Member Up
                0060b0-df2a00   Newstack-2      Member Up
Others:         001083-3c09c0   DEFAULT_CONFIG  Candidate
                0060b0-e94300   DEFAULT_CONFIG  Candidate
                080009-918f80   DEFAULT_CONFIG  Candidate

Actions->  Back  Next page  Prev page  Help

Return to previous screen.
Use up/down arrow keys to scroll to other entries, left/right arrow keys to
change action selection, and <Enter> to execute action.
```

For status descriptions, see the table on page 83.

This column lists the MAC Addresses for switches discovered (in the local subnet) that are configured for Stacking.

If you know the MAC addresses for these Members, you can move them between stacks in the same subnet.

Figure 37. Example of How the Stacking Status (All) Screen Helps You Find Member MAC Addresses

3. In the Stacking Status (All) screen, find the Member switch that you want to move and note its MAC address, then press [B] (for **Back**) to return to the Stacking Menu.
4. Display the Commander's Stack Management screen by selecting

4. Stack Management

(For an example of this screen, see figure 34 on page 72.)

5. Press [A] (for **Add**) to add the Member. You will then see a screen listing any available candidates. (See figure 35 on page 73.) Note that you will not see the Member you want to add because it is a Member of another stack and not a Candidate.)
6. Either accept the displayed switch number or enter another available number. (The range is 0 - 15, with 0 reserved for the Commander.)
7. Use the downarrow key to move the cursor to the MAC Address field, then type the MAC address of the desired Member you want to move from another stack.
8. Do one of the following:
 - If the stack containing the Member you are moving has a Manager password, press the downarrow key to select the Candidate Password field, then type the password.
 - If the stack containing the Member you want to move does not have a password, go to step 9.
9. Press [Enter] to return to the Actions line, then press [S] (for **Save**) to complete the Add process for the selected Member. You will then see a screen similar to the one in figure 34 on page 72, with the newly added Member listed.

Note: If the message **Unable to add stack member: Invalid Password** appears in the console menu's Help line, then you either omitted the Manager password for the stack containing the Member or incorrectly entered the Manager password.

Note: You can move a Member from one stack to another by entering the MAC address for the destination stack Commander in the Member's **Commander MAC Address** field. Using this method moves the Member to another stack without a need for knowing the Manager password in that stack, but also blocks access to the Member from the original Commander.

Removing a Member from a Stack

These rules affect removals from a stack:

- When a Candidate becomes a Member, its **Auto Join** parameter is automatically set to **No**. This prevents the switch from automatically rejoining a stack as soon as you remove it from the stack.
- When you use the Commander to remove a switch from a stack, the switch rejoins the Candidate pool for your subnet (broadcast domain), with **Auto Join** set to **No**.
- When you remove a Member from a stack, its switch number (**SN**) becomes available for assignment to another switch that you may subsequently add to the stack. The default switch number used for an add is the lowest unassigned number in the Member range (1 - 15; 0 is reserved for the Commander).

To remove a Member from a stack, use the Stack Management screen.

1. From the Main Menu, select:

8. Stacking...

4. Stack Management

You will then see the Stack Management screen:

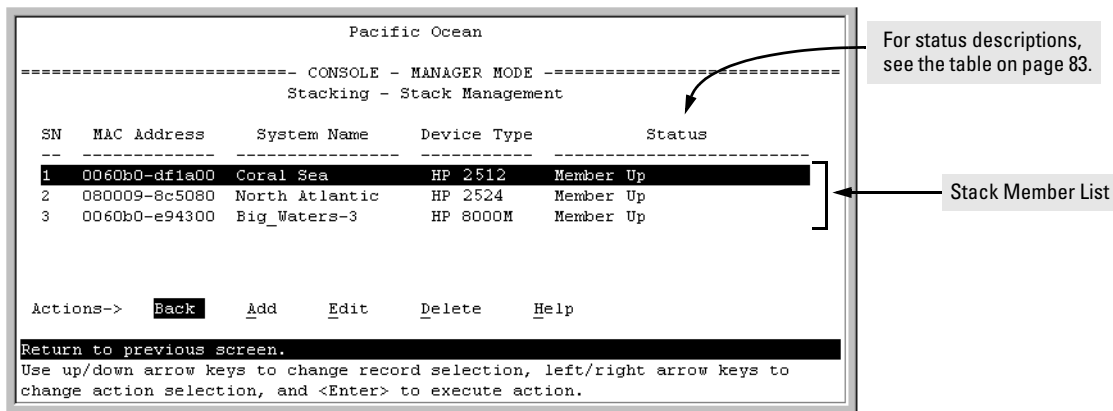


Figure 38. Example of Stack Management Screen with Stack Members Listed

2. Use the downarrow key to select the Member you want to remove from the stack.




Figure 39. Example of Selecting a Member for Removal from the Stack

3. Type [D] (for **Delete**) to remove the selected Member from the stack. You will then see the following prompt:

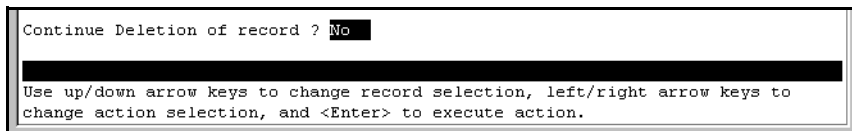


Figure 40. The Prompt for Completing the Deletion of a Member from the Stack

4. To continue deleting the selected Member, press the Space bar once to select **Yes** for the prompt, then press [Enter] to complete the deletion. The Stack Management screen updates to show the new stack Member list.

Accessing Member Switches To Make Configuration Changes and Monitor Traffic

After a Candidate becomes a Member, you can access its console interface for the same configuration and monitoring that you would do through a Telnet or direct-connect access.

1. From the Main Menu, select:

8. Stacking...

5. Stack Access

You will then see the Stack Access screen:

```
Pacific Ocean
-----
CONSOLE - MANAGER MODE -----
Stacking - Stack Access

SN      MAC Address      System Name      Device Type      Status
-----
0       0060b0-880a80    Pacific Ocean    HP 4108          Commander Up
1       0060b0-df1a00    Coral Sea        HP 2524          Member Up
2       080009-8c5080    North Atlantic   HP 8000M         Member Up

Actions->  Cancel      eXecute      Help

Return to previous screen.
Use arrow keys to change field selection
```

For status descriptions, see the table on page 83.

Figure 41. Example of the Stack Access Screen

Use the downarrow key to select the stack Member you want to access, then press [X] (for **eXecute**) to display the console interface for the selected Member. For example, if you selected switch number 1 (system name: Coral Sea) in figure 41 and then pressed [X], you would see the Main Menu for the switch named Coral Sea.

```
Coral Sea
-----
TELNET - MANAGER MODE -----
Main Menu

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

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

Main Menu for stack Member named "Coral Sea" (SN = 1 from figure 41)

Figure 42. The eXecute Command Displays the Console Main Menu for the Selected Stack Member

2. You can now make configuration changes and/or view status data in the selected Member in the same way that you would if you were directly connected or telnetting into the switch.
3. When you are finished accessing the selected Member, do the following to return to the Commander's Stack Access screen:
 - a. Return to the Member's Main Menu
 - b. Press [0] (for Logout), then [Y] (for Yes).
 - c. Press [Return].

You should now see the Commander's Stack Access screen. (For an example, see figure 41 on 77.)

Monitoring Stack Status

Using the stacking options in the console interface for the Commander, you can view stacking data for the Commander, any Member, any Candidate, or all stacking commanders, Members, and Candidates in the subnet (broadcast domain). (If you are using VLANs in your stack environment, see "Stacking Operation with a Tagged VLAN" on page 83.) This can help you in such ways as determining the stacking configuration for individual switches, identifying stack Members and Candidates, and determining the status of individual switches in a stack.

Screen Name	Commander	Member	Candidate
Stack Status (This Switch)	<ul style="list-style-type: none"> • Commander's stacking configuration • Data on stack Members: <ul style="list-style-type: none"> – Switch Number – MAC Address – System Name – Device Type – Status 	<ul style="list-style-type: none"> • Member's stacking configuration • Member Status • Data identifying Member's Commander: <ul style="list-style-type: none"> – Commander Status – Commander IP Address – Commander MAC Address 	Candidate's stacking configuration
Stack Status (All)	<p>Lists devices by stack name or Candidate status (if device is not a stack Member). Includes:</p> <ul style="list-style-type: none"> • Stack Name • MAC Address • System Name • Status 	Same as for Commander.	Same as for Commander.

Status for All Switches with Stacking Enabled. This procedure displays the general status of all switches in the subnet (broadcast domain) that have stacking enabled.

1. Go to the console Main Menu for any switch configured for stacking and select:

8. Stacking ...

2. Stacking Status (All)

You will then see a Stacking Status screen similar to the following:

```

Pacific Ocean
-----
-- CONSOLE - MANAGER MODE -----
Stacking - Stacking Status (All)
-----
Stack Name      MAC Address    System Name    Status
-----
Big Waters      0060b0-880a80 Pacific Ocean  Commander Up
                0060b0-df1a00 Coral Sea     Member Up
Newstack        080009-8c5080 North Atlantic Member Up
                001083-c3fc00 Newstack-0    Commander Up
                080009-918f80 Newstack-1    Member Up
                0060b0-df2a00 Newstack-2    Member Up
Others:         001083-3c09c0 DEFAULT_CONFIG Candidate
                0060b0-e94300 DEFAULT_CONFIG Candidate
                080009-918f80 DEFAULT_CONFIG Candidate

Actions->  Back  Next page  Prev page  Help
-----
Return to previous screen.
Use up/down arrow keys to scroll to other entries, left/right arrow keys to
change action selection, and <Enter> to execute action.
  
```

Figure 43. Example of Stacking Status for All Detected Switches Configured for Stacking

Commander Status. To display the status for a Commander, go to the console Main Menu for the switch and select:

8. Stacking ...

1. Stacking Status (This Switch)

You will then see the Commander's Stacking Status screen:

```
Pacific Ocean
----- CONSOLE - MANAGER MODE -----
                Stacking - Stacking Status (This Switch)

Stack State      : Commander
Transmission Interval : 60
Stack Name       : Big_Waters Number of members      : 2
Auto Grab        : No           Members unreachable   : 0

SN   MAC Address   System Name   Device Type   Status
-----
0    0060b0-880a80 Pacific Ocean HP 4108      Commander Up
1    0060b0-dfia00 Coral Sea    HP 2524      Member Up
2    080009-8c5080 North Atlantic HP 8000M     Member Up

Actions->  Back  Help
Return to previous screen.
Use arrow keys to change action selection and <Enter> to execute action.
```

Figure 44. Example of the Commander's Stacking Status Screen

Member Status. To display the status for a Member:

1. Go to the console Main Menu for the Commander switch and select

8. Stacking ...

5. Stack Access

2. Use the downarrow key to select the Member switch whose status you want to view, then press [X] (for **eXecute**). You will then see the Main Menu for the selected Member switch.
3. In the Member's Main Menu screen, select

8. Stacking ...

1. Stacking Status (This Switch)

You will then see the Member's Stacking Status screen:


```

Coral Sea
----- TELNET - MANAGER MODE -----
Stacking - Stacking Status (This Switch)

Stack State      : Member
Transmission Interval : 60
Switch Number    : 1
Stack Name       : Big_Waters
Member Status    : Joined Successfully
Commander Status  : Commander Up
Commander IP Address : 13.28.227.102
Commander MAC Address : 0060b0-880a80

Actions->  Back  Help

Return to previous screen.
Use arrow keys to change action selection and <Enter> to execute action.

```

Figure 45. Example of a Member's Stacking Status Screen

Candidate Status. To display the status for a Candidate:

1. Use Telnet or a direct serial port connection to access the console Main Menu for the Candidate switch and select

8. Stacking ...

1. Stacking Status (This Switch)

You will then see the Candidate's Stacking Status screen:

```

Coral Sea
----- TELNET - MANAGER MODE -----
Stacking - Stacking Status (This Switch)

Stack State      : Candidate
Transmission Interval : 60
Auto Join        : No

Actions->  Back  Help

Return to previous screen.
Use arrow keys to change action selection and <Enter> to execute action.

```

Figure 46. Example of a Candidate's Stacking Screen

SNMP Community Operation in a Stack

Community Membership. When a Candidate becomes a Member of a stack, it automatically becomes a Member of any SNMP community to which the Commander belongs, even though any community names configured in the Commander are not propagated to the Member's SNMP Communities screen. However, if a Member has its own (optional) IP addressing, it can belong to SNMP communities to which other switches in the stack, including the Commander, do not belong. For example:

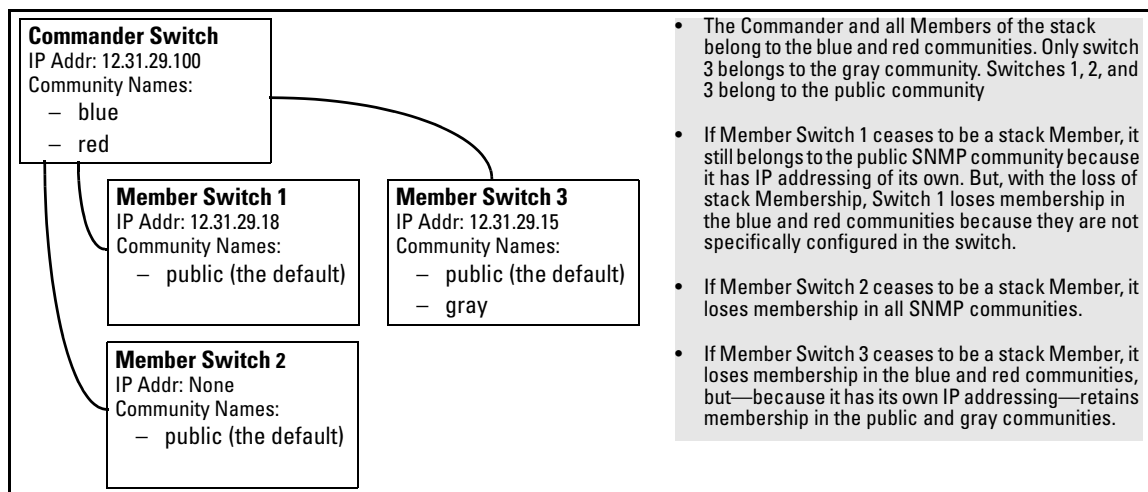


Figure 47. Example of SNMP Community Operation with Stacking

SNMP Management Station Access to Members Via the Commander. To use a management station for SNMP Get or Set access through the Commander's IP address to a Member, you must append `@sw<switch number>` to the community name. For example, in figure 47, you would use the following command in your management station to access switch 1's MIB using the blue community:

```
snmpget <MIB variable> 12.31.29.100 blue@sw1
```

Note that because the gray community is only on switch 3, you could not use the Commander IP address for gray community access from the management station. Instead, you would access switch 3 directly using the switch's own IP address. For example:

```
snmpget <MIB variable> 12.31.29.15 gray
```

Note that in the above example (figure 47) you cannot use the public community through the Commander to access any of the Member switches. For example, you can use the public community to access the MIB in switches 1 and 3 by using their unique IP addresses. However, you must use the red or blue community to access the MIB for switch 2.

```
snmpget <MIB variable> 12.31.29.100 blue@sw2
```

Stacking Operation with a Tagged VLAN

To use a tagged VLAN in a stacking environment, the following criteria applies:

- For each switch in the stack, the tagged VLAN must be the default VLAN on links used for stacking. In the console interface, the default VLAN is the first VLAN listed in the VLAN Names and VLAN Port Assignment screens. The web browser interface attaches a **"Default"** label to the VLAN ID for the default VLAN.
- The VLAN ID (VID) for the (tagged) default VLAN must be the same for all switches in the stack. (The default VLAN name can differ among switches in the stack.)

Status Messages

Stacking screens display these status messages:

Message	Condition	Action or Remedy
Candidate Auto-join	Indicates a switch configured with Stack State set to Candidate, Auto Join set to Yes (the default), and no Manager password.	None required
Candidate	Candidate cannot automatically join the stack because one or both of the following conditions apply: <ul style="list-style-type: none">• Candidate has Auto Join set to No.• Candidate has a Manager password.	Manually add the candidate to the stack.
Commander Down	Member has lost connectivity to its Commander.	Check connectivity between the Commander and the Member.
Commander Up	The Member has stacking connectivity with the Commander.	None required.
Mismatch	This may be a temporary condition while a Candidate is trying to join a stack. If the Candidate does not join, then stack configuration is inconsistent.	Initially, wait for an update. If condition persists, reconfigure the Commander or the Member.
Member Down	A Member has become detached from the stack. A possible cause is an interruption to the link between the Member and the Commander.	Check the connectivity between the Commander and the Member.
Member Up	The Commander has stacking connectivity to the Member.	None required.
Rejected	The Candidate has failed to be added to the stack.	The candidate may have a password. In this case, manually add the candidate. Otherwise, the stack may already be full. A stack can hold up to 15 Members (plus the Commander).

Changes to the Web Browser Interface for Commander Switches

On a Commander, the web browser interface includes the Stack Access pull-down menu and buttons shown below. For information on how to use the web browser interface to configure and manage stacking, see the online Help for the web browser interface.

These three buttons appear only in the web browser interface for a switch configured as a Commander. For more information, see the online Help provided for the web browser interface.

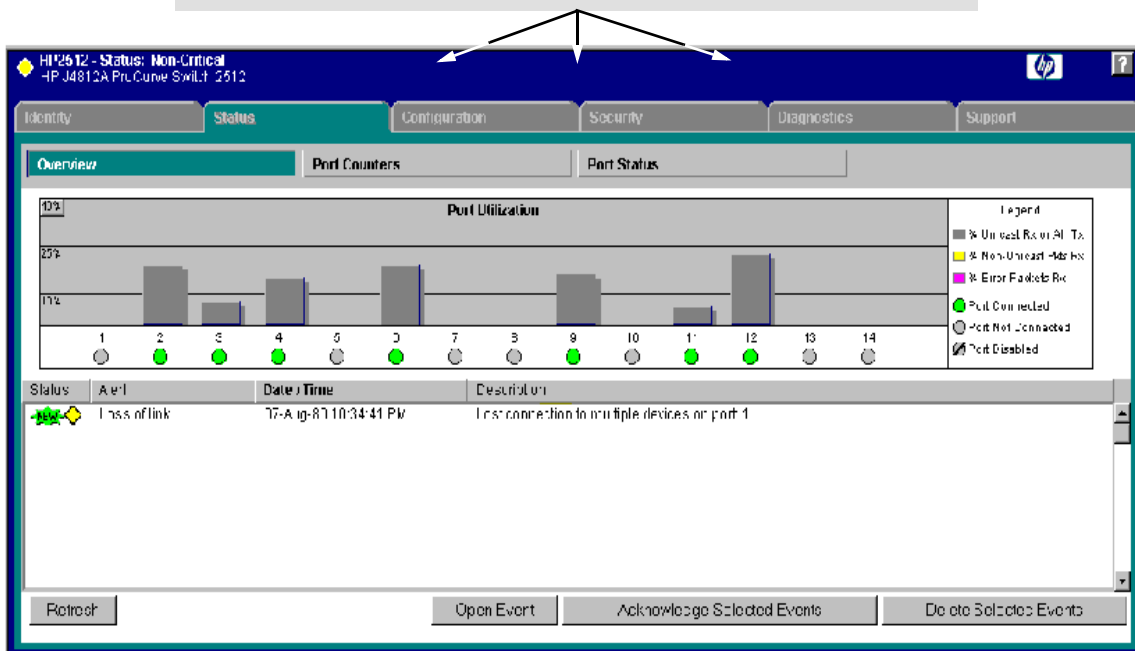


Figure 48. Example of a Web Browser Interface for a Switch Configured as a Commander

Using the Auto-10 Port Configuration Option

Auto-10 is a new port mode that allows the port to negotiate between half-duplex (Hdx) and full-duplex (Fdx) while keeping speed at 10 Mbps. HP recommends Auto-10 for links between 10/100 autosensing ports connected with Cat 3 cabling. (Cat 5 cabling is recommended for 100 Mbps links.)

```
DEFAULT CONFIG
-----
----- CONSOLE - MANAGER MODE -----
Switch Configuration - Port Settings

Port   Type      Enabled  Mode      Flow Ctrl  Beast Limit
-----+-----
A1     10/100TX  | Yes     Auto-10   Disable    0
A2     10/100TX  | Yes     Auto      Disable    0
A3     10/100TX  | Yes     Auto      Disable    0
```

The Auto-10 Mode option is available in the Port Settings screen of the console interface.

In the web browser interface, configure Auto-10 in the Port Configuration window. (Click on the **Configuration** tab, then the [Port Configuration] button.)

Updates to VLAN Configuration Options

VLAN Enhancement to the Web Browser Interface. You can now enable or disable VLAN support, and add, remove, and rename VLANs in the web browser interface. Click on Configuration, then VLAN Configuration. (To change the number of VLANs allowed, you must still use the console interface.)

Assigning Untagged VLAN Status to a Port in a Multiple VLAN Environment. When multiple VLANs exist on a switch, only one VLAN can be untagged for each port. When you first add a VLAN to a switch, the default setting on that VLAN is **No** for all ports. On the web browser interface, if you subsequently reconfigure a port to **Untagged** for the new VLAN while there is an **Untagged** setting on another VLAN for the same port, the switch automatically reconfigures the other VLAN setting to **No**. For example, if you configure port A1 as **Untagged** for **2nd_VLAN**, then the switch automatically reconfigures **DEFAULT_VLAN** for port A1 as **No**. For more information, see the online Help provided with the web browser interface.

Enhanced Multimedia Traffic Filtering

With IP Multicast (IGMP) enabled, ports listen to ("snoop") both IGMP messages and IP multicast streams. Thus, the switch learns the multicast traffic stream MAC addresses and begins filtering the stream right away. As soon as a "join" is heard from a client, the switch begins forwarding the requested IP multicast traffic out the port from which the join was heard. Multicast traffic destined for "reserved" IP multicast addresses is flooded out all ports.

FAQs from the HP ProCurve Web Site

This section provides answers to frequently asked questions regarding HP ProCurve Switch 1600M, 2400M, 2424M, 4000M, and 8000M operation.

Auto-Negotiation	Below
Cabling	88
Features	88
Gigabit Stacking	91
Modules	91
Spanning Tree Protocol	93
Troubleshooting	95
Trunking (HP and Fast EtherChannel—)	95


Auto-Negotiation

Q: Does the factory configure the 10/100 auto-sensing modules with auto-sensing enabled?

A: Yes.

Q: Is 10/100Mbps auto-negotiation the same as Plug-n-Play?

A: No. The following configuration will cause severe network problems:



The hub, switch, or router will correctly sense (not auto-negotiate) the 10Mbps or 100Mbps speed. Since the end node was configured for a specific speed and duplex state, and therefore does not negotiate, the hub, switch, or router will choose the communication mode specified by the 802.3u standard, namely half-duplex.

With one device running at half-duplex and the device on the other end of the connection at full-duplex, the connection will work reasonably well at low levels of traffic. At high levels of traffic the full-duplex device (end node, in this case) will experience an abnormally high level of CRC or alignment errors. The end users usually describe this situation as, "Performance seems to be approximately 1Mbps!." Often, end nodes will drop connections to their servers.

In this same situation, the half-duplex device will experience an abnormally high level of late collisions.

The network administrator must take care to verify the configuration of each network device during installation. Also, check the operational mode of each network device. That is, check both how you configured it and also that it comes up as you expect, for example, at 10Mbps/half-duplex.

Q: Is Gigabit Ethernet auto-negotiation the same as Plug-n-Play?

A: No. By the time the IEEE issued the 802.3z specification, they knew about the 10/100Mbps auto-negotiation problem (see the FAQ "Is 10/100Mbps auto-negotiation the same as Plug-n-Play?"). To prevent it, 802.3z auto-negotiation requires that, if one side of a connection is configured to auto-negotiate, the other side must also auto-negotiate if the connection is to come up. In other words, if a switch is configured to auto-negotiate and its attached end node is configured to, say, 1000Mbps/full-duplex, the 803.2z spec requires that the switch NOT allow the link to come up.

Cabling

Q: Can I use category 3 cables with the HP ProCurve Switch 100/1000Base-T Module?

A: No, only category 5 100-ohm UTP or STP cables are supported. In fact, for the most robust connections you should use cabling that complies with the Category 5E specifications, as described in Addendum 5 to the TIA-568-A standard (ANSI/TIA/EIA-568-A-5).

Q: What is the maximum length for cables used with the HP ProCurve Switch 100/1000Base-T Module?

A: The maximum length is 100 meters using category 5, 100-ohm UTP or STP cable. This distance is correct for the IEEE 802.3ab specification.

Features


Q: What are the differences between the HP ProCurve Switch 2424M (product number J4093A) and the HP ProCurve Switch 2400M (product number J4122A)?

A: The HP ProCurve Switch 2424M has 24 10/100Base-TX autosensing ports like the HP ProCurve Switch 2400M, but the HP ProCurve Switch 2424M also has a module slot that can be used for stacking, gigabit connectivity or port expansion.

Q: What would a topology look like that has a maximum meshed switch hop count of 5?

A: Up to 12 switches are supported in a switch mesh domain (See the Management and Configuration Guide), and a maximum meshed switch hop count of five is allowed in the path connecting two nodes via a switch mesh domain topology.

"Hops" refers to inter-switch links, not the number of switches crossed. The limit is 5 meshed links. The rest of this discussion will use the word links instead of hops. In the diagram below, the numbers and arrows above the switches show the link count value in the mesh protocol packets. There are six switches between End Node A and End Node B, but only five links:




The switches learn the mesh topology by sending mesh protocol packets to each other. Each switch initiates this learning process by sending a mesh protocol packet with a link count of zero. When its neighbor switch receives this packet, the neighbor increments the link count (for example, from zero to one), then propagates the packet to the next switch. Switches do not send mesh protocol packets with a link count of 5 or larger.

In the diagram above, you can see that Switch 1 sends a mesh protocol packet containing a link count of zero. Switch 2 receives it and sends a mesh protocol packet with a link count of one. Switch 3 receives it and sends a mesh protocol packet with a link count of two. And so on.

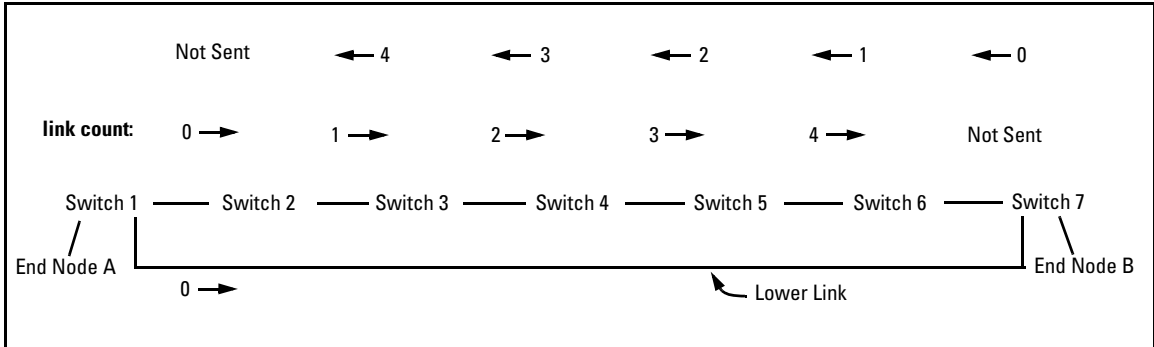
Note that if we were to violate the mesh link count limit by adding a 7th switch between Switch 6 and End Node B, then Switch 7 would not receive any mesh protocol packets from Switch 1. Since Switch 7 would be unaware of the Switch 1, end nodes on Switch 1 would not be able to communicate with end nodes on Switch 7.

Consider now an illegal topology:



Switch 7 does not receive mesh protocol packets from Switch 1.

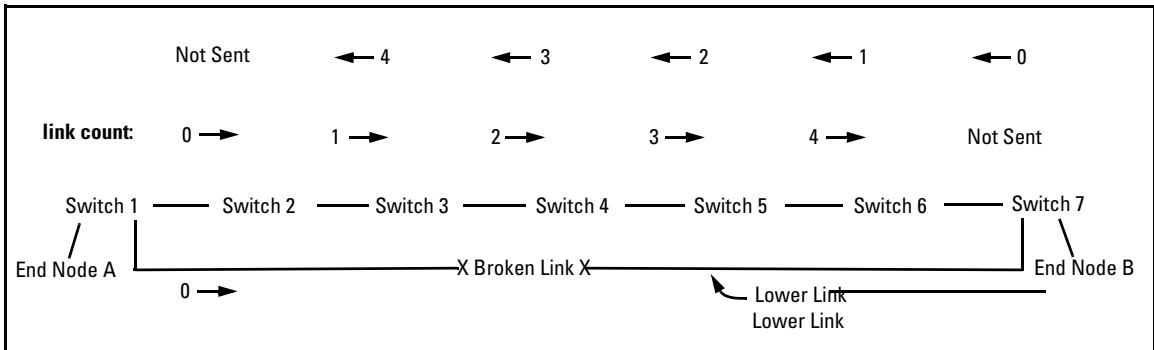
If we add an apparently redundant link:



In this topology, we appear to have redundant connections. Switch 1 is connected to Switch 7 through the upper link, and again through the "lower" link. You can see that Switch 7 learns that it is directly connected to Switch 1 through the "lower" link.

Switch 1 and Switch 7 do not know that they have a connection through the upper link, since mesh protocol packets will not travel across that many links. All traffic between End Node A and End Node B will travel along the lower link. In other words, we do not have a usable redundant link, since the topology is illegal.

In the diagram above, if the lower link becomes broken, we get the situation shown below:



In the above topology, Switch 7 knows that it no longer receives mesh protocol packets from Switch 1 (through the lower link), but it does not know that it has an alternate path through the upper link. Connectivity between End Node A and End Node B is lost, and the mesh protocol cannot recover it.

Conclusion: For the switch mesh to work properly you must satisfy both limits:

- A maximum of 5 meshed switch links between nodes; and
- A maximum of 12 switches per mesh.

Gigabit Stacking

Q: Does the HP ProCurve Switch 2424M support stacking?

A: The HP ProCurve Switch 2424M has a slot for the two port HP ProCurve Switch 2424M Gigabit Stacking Module (product number J4130A). This module stacks up to 7 switches with the HP ProCurve Switch 2424M Gigabit Stacking Kit (product number J4116A). One stacking module must be purchased for each switch and one stacking kit must be purchased for every stack of two switches. The stacking module can also be used as a dual gigabit uplink module with the HP ProCurve Gigabit-SX Transceiver (product number J4131A) and/or the HP ProCurve Gigabit-LX Transceiver (product number J4132A). In addition, the following modules, which are supported in the HP ProCurve Switch 8000M, 4000M, 2400M, and 1600M, are also supported in the HP ProCurve Switch 2424M for uplink connectivity or for additional desktop ports:

- HP ProCurve Switch 100/1000Base-T Module (product number J4115A and J4115B)
- HP ProCurve Switch 10/100Base-T Module (product number J4111A)
- HP ProCurve Switch 100Base-FX Module (product number J4112A)
- HP ProCurve Switch Gigabit-SX Module (product number J4113A)
- HP ProCurve Switch Gigabit-LX Module (product number J4114A)
- HP ProCurve Switch 10Base-FL Module (product number J4118A)

Q: Can the HP ProCurve Switch 2424M Gigabit Stacking Module (product number J4130A) be used in the HP ProCurve Switch 8000M, 4000M, 2400M or 1600M?

A: No, this stacking module was designed specifically for the HP ProCurve Switch 2424M. It will not slide into a slot on any other switch. But the optional modules supported on the HP ProCurve Switch 8000M, 4000M, 2400M, and 1600M are supported on the HP ProCurve Switch 2424M.

Q: Does the HP ProCurve Switch 2424M Gigabit Stacking Module support trunking?

A: Yes, the HP ProCurve Switch 2424M Gigabit Stacking Module supports trunking. Port trunk links can be configured only between like media types:

- Gigabit-SX with Gigabit-SX
- Gigabit-LX with Gigabit-LX
- HP ProCurve Switch 2424M Gigabit Stacking Kit transceiver with HP ProCurve Switch 2424M Gigabit Stacking Kit transceiver

Modules

See also “Gigabit Stacking” on page 91 and “Cabling” on page 88.

Q: Does the 100Base-FX module support 10Base-FL operation?

A: No, there is a separate card to support 10Base-FL fiber optic Ethernet connectivity.

Q: What is the supported distance of the Gigabit-SX module?

A:

Model Bandwidth	Distance Covered if 62.5 Core Diameter	Distance Covered if 50 Core Diameter
160 MHz.km	220 meters	N/A
200 MHz.km	275 meters	N/A
400 MHz.km	N/A	500 meters
500 MHz.km	N/A	550 meters

The distances noted above are correct for the IEEE Draft P802.3z/D4.2 (Gigabit Draft 4.2) specification.

The better the quality of the cable, the greater the modal bandwidth and the greater the distance supported.

Q: What are the features of the HP ProCurve Switch 100/1000Base-T Module?

A: The HP ProCurve Switch 100/1000Base-T Module (J4115A and J4115B) is an accessory that can be added to the HP ProCurve Switch 1600M, 2424M, 4000M, and 8000M. The HP ProCurve Switch 100/1000Base-T Module provides one twisted-pair port with an RJ-45 connector for 100Mbps or 1000Mbps (Gigabit) operation over category 5 cable. Like all of the switch modules available for this line of products, the HP ProCurve Switch 100/1000Base-T Module is shipped ready for network operation as soon as a viable network cable is connected. The LEDs on this module provide information on the link status, network activity, connection bandwidth, and communication mode (half or full duplex). And, as with all other HP ProCurve modules, it features "hot-swap" operation, meaning addition or replacement of a module can be achieved without having to shut down the switch (changing the module type in a given slot does require a switch reset).

Q: Do I need a software update on my switch to use the HP ProCurve Switch 100/1000Base-T Module?

A: Yes. Your HP ProCurve Switch requires software version C.07.20 or later to support the J4115A HP ProCurve Switch 100/1000Base-T Module, and software version C.07.27 or later to support the J4115B HP ProCurve Switch 100/1000Base-T Module. The new software is found on the floppy disk included with the module. You can also download the software from the HP ProCurve Networking web site (select "Free Software Updates" on the left side of the home page). After installation of the software you can install the module into a slot on the switch.

Q: Is the HP ProCurve Switch 100/1000Base-T Module pre-configured with auto-sensing enabled?

A: Yes. In the default state, 100/1000Mbps connections can auto-negotiate for MDI/MDI-X, Full/Half Duplex, 100/1000Mbps, Flow control on/off, and, in 1000Mbps only, for master/slave.

Q: What configuration options are available on the HP ProCurve Switch 100/1000Base-T Module?

A: Three options are available:

- Auto - The module auto negotiates connection speed (100 or 1000Mbps), communication mode (full or half duplex), and MDI or MDI-X port operation.
- 100 Full - The module is forced to 100Mbps speed and the communication mode is set to full duplex.
- 100 Half - The module is forced to 100Mbps speed and the communication mode is set to half duplex.

Note

If you configure the port to one of the fixed 100Mbps modes, the port will then operate only as an MDI-X port.

Spanning Tree Protocol

See also “Troubleshooting” on page 95.

Q: What should I do if my switch (which has switched ports to the desktop) reports many Spanning Tree Protocol (STP) topology changes, even if I am not having connectivity problems?

A: You can view how many topology changes are occurring by looking at the Topology Change Count in the Status and Counters—Spanning Tree Information screen on your HP ProCurve Switch 8000M, 4000M, 2424M, 2400M, or 1600M.

STP was developed to manage switch-to-switch links, or bridge-to-bridge links as they were initially called. For most of STP's lifetime, a topology change was a serious event as it indicated a change in the LAN topology. It also indicated a likelihood that STP had to bring (at least) a portion of the LAN down for awhile in order to rebuild a new working topology. Of course, in a LAN we expect the topology to change only rarely.

Over the past few years, the cost of switch ports has dropped dramatically. Users can now afford to dedicate switch ports to end nodes.

Once an STP topology is stable, establishing an Ethernet link on a switch port results in a topology change. This happens every time a user powers up their PC (assuming that the PC is directly connected to a switch port), resets the PC, or brings up the PC's network stack. The establishment of this link causes the topology change count to increment in:

1. the switch to which the end node is directly attached; and
2. "upstream" switches. That is, the root switch and the switches between the root switch and the switch to which the end node is directly connected.

This scenario does not result in any topology issues or changes. It does not result in any loss of connectivity in the LAN. If the incrementing of your topology change count is due to this type of scenario, you do not need to take any action.

The above discussion applies to Hewlett-Packard switches in general and HP ProCurve switches in particular when configured in STP Normal mode. When an HP ProCurve switch's port is configured in STP Fast mode, the switch will not increment the topology change count as a result of link changes on that port. Please see the switch's Management and Configuration Guide for details on normal and fast modes.

Q: When I power on my PC, I get the message "a file server could not be found." How do I fix that?

A: This is a well-known issue given the following situation:

1. The PC is directly-connected to a switch.
2. The PC is running Novell's VLMs or Client32.
3. The switch has Spanning Tree Protocol (STP) enabled.

In this situation, when the directly-connected PC is powered on, the switch senses linkbeat on that port. This causes the switch to go through the four Spanning Tree states: blocking, listening, learning, and forwarding. It takes 30 seconds for the switch to complete that sequence and begin forwarding packets to and from that port. During those 30 seconds, Novell sends 3 requests for a server, then stops looking. By the time Spanning Tree completes its job, Novell reports that "a file server could not be found."

There are several workarounds available:

1. Disable Spanning Tree on the switch (if Spanning Tree is not needed, i.e. no loops in the network topology).
2. For VLMs, add a "pause" just after calling VLM.EXE in STARTNET.BAT. When the user reboots a PC, have them wait at least 30 seconds before continuing the sequence. This workaround is documented on Novell's Knowledgebase (www.support.novell.com, search for document 2920460).
3. For Client32, add a registry entry in the PC, as documented on Novell's Knowledgebase (search for document 2925582).

In the Fall of 1998, HP released switch firmware (version C.05.07 or greater) with an enhancement to resolve this timing problem between Novell and STP. The enhancement allows users to configure Spanning Tree so that it does not go through the 4 states, on a port-by-port basis. Instead, for those configured ports, Spanning Tree will immediately begin forwarding packets to and from the port. This allows Novell clients to communicate with the server as soon as the network card (NIC) is enabled. After that, the switch continues to listen for and send Spanning Tree packets on those

configured ports. This protects the user who might inadvertently connect a hub or switch to that port and create a network loop—Spanning Tree will detect the loop after a short time, since the port listens for and sends STP packets on that port.

Troubleshooting

Q: Why won't my iMac, which is directly attached to an HP ProCurve switch, NetBoot from an OS X Server, even though it worked correctly when my iMac was plugged directly into a hub?

A: This issue is not specific to HP switches. Rather, it is a timing problem in early versions of the iMac Boot ROM. Apple Computer, Inc has fixed this NetBoot issue by releasing both an updated Boot ROM image and Mac OS X Server version 1.2. Please advise customers with iMac NetBoot issues to contact Apple Computer, Inc Technical Support.

Q: Why is my Macintosh system unable to use AppleTalk services?

A: Possible symptoms include: no AppleTalk services, only local network AppleTalk services, performance problems, and intermittent network services. If you remove the Macintosh from its dedicated switch (or routing switch or router) port and connect it to a hub, the problem goes away.

If the switch (or routing switch or router) has Spanning Tree Protocol enabled, see Apple Computer, Inc's Tech Info Library entry "Spanning Tree Protocol: AppleTalk Issues" at <http://til.info.apple.com/techinfo.nsf/artnum/n30922>

Trunking (HP and Fast EtherChannel—FEC)

See also "Gigabit Stacking" on page 91.

Q: Can the HP ProCurve Switch 8000M, 4000M, 2424M, 2400M, or 1600M connect to the HP AdvanceStack Switch 2000 or 800T with HP Port Trunking?

A: Yes. The HP Port Trunking implemented in the HP ProCurve Switch 8000M, 4000M, 2424M, 2400M, and 1600M has been enhanced from the version implemented in the HP AdvanceStack Switch 2000 and 800T. But you can still connect any of these switches together using HP Port Trunking.

Q: Is the Fast EtherChannel feature compatible with Cisco internetworking equipment?

A: Yes.

Q: How many links can I connect to my Unix server?

A: You can configure anywhere from 1-4 100Base-T links using the Fast EtherChannel feature.


VLANs

Q: How many VLANs can be configured on the HP ProCurve Switch 8000M, 4000M, 2424M, 2400M, and 1600M?

A: These switches support up to 30 VLANs when using software version C.06.06 or later.

Q: What is the recommended way to connect multiple VLANs between a routing switch and a layer 2 switch?

A: The diagram below illustrates the question.



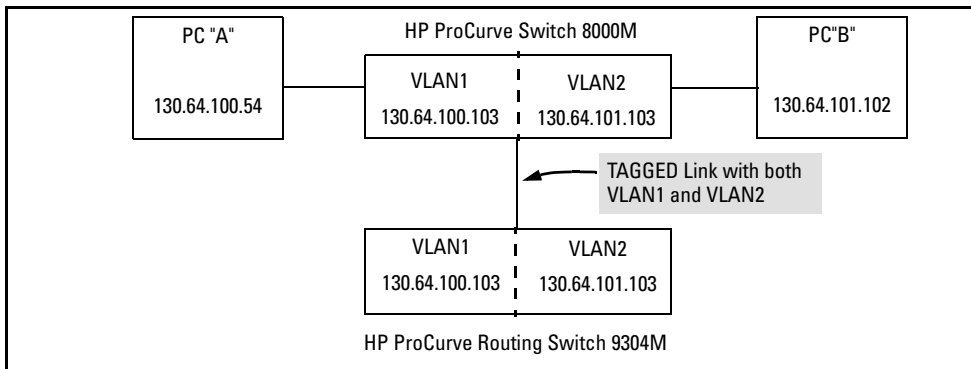
The following HP switches provide VLANs and have a single MAC/Ethernet address (filtering) table: Switch 800T, 2000, 1600M, 2400M, 2424M, 4000M, 8000M. In the diagram above we show a Switch 8000M, but the following discussion applies to all of the switches listed in the previous sentence. The HP ProCurve Routing Switch 9304M, 9308M, or 6308M-SX, as a default gateway, has a single MAC address (for all of its VLANs) if using virtual Ethernet interfaces. In the diagram above we show a 9304M, but this could be a 9308M or 6308M-SX as well.

Let's consider PC "A" attempting to send an IP packet to PC "B". PC "A" will send the 8000M a packet with the 9304M's MAC address in the destination field. If the 8000M has not yet learned this MAC address, the 8000M will flood the packet out all of its VLAN1 ports, including the VLAN1 link to the 9304M. The 9304M will then route the packet toward PC "B" via its link with the 8000M's VLAN2 connection. The 8000M will enter the 9304M's MAC address into its MAC address table as located in VLAN2. The 8000M will also forward the packet to PC "B".

Let's consider a second packet that PC "A" sends to PC "B". PC "A" sends the packet, again addressed to the 9304M's MAC address, to the 8000M. The 8000M will check its address table and find that the 9304M appears to be located on VLAN2. Since the 8000M believes that this MAC address is not located on VLAN1, the switch will discard the packet.

Later, when the 9304M transmits a packet to the 8000M via the VLAN1 link, the 8000M will update its address table to indicate that the 9304M's MAC address is located in VLAN1 instead of VLAN2. As you can see, the 8000M's location information for the 9304M's MAC address will vary over time between VLAN1 and VLAN2. For this reason, some packets directed through the 8000M for the 9304M's MAC address will be discarded. Performance may appear to be poor or connectivity may appear to be broken.

To avoid this issue, simply use one cable between the 8000M and the 9304M instead of two, making sure that the two VLANs use tags on that link, as shown below.



(This page intentionally left blank)

Release History

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

Release C.09.22

Problems Resolved in Release C.09.22

- **Linktest (PR_1000001639)** — The switch will respond to “linktest” packets that have a broadcast or multicast MAC destination address.
- **Meshing (PR_92500)** — CDP packets will not forward properly through a mesh when CDP is disabled, resulting in a CDP packet loop.
- **Meshing (PR_93999)** — A meshed switch may experience a loop during very high traffic levels and large amounts of port toggling (Link Up, Down, Up,...).
- **Security (PR_87930/Switch 4520)** — The switch will allow a user set up as ‘Operator’ the same privileges as ‘Manager’.
- **Task Manager (PR_1000002122)** — The switch may suspend an arbitrary task (e.g. TELNET, STP, etc.) in the event a memory corruption occurs.
- **VLAN (PR_1000004334)** — If VLANs are not enabled, then at the time the user enables VLAN monitoring only ports with Link will pass traffic. Ports that later get Link will not pass traffic. Statistics will report (good) transmits and receives even for ports that are not passing traffic.
- **Web Interface (PR_88745)** — A user may not be able to add a new VLAN after a VLAN is deleted.
- **Web Interface (PR_92078)** — Page may not fully load after making changes under the Device Features tab.
- **Web Interface (PR_98500)** — With Sun Java 1.3.1, the browser window may spontaneously close.
- **Web Interface (PR_98919)** — With Java 1.4.2.1, "clear changes" button doesn't clear STP.
- **Web Interface (PR_1000000449)** — With Sun java 1.3.1/1.4.1, nothing will display when the users clicks on the VLAN Configuration button.
- **Web Interface (PR_1000000452)** — Resetting the switch leads to the URL aol.co.uk.

Release C.09.19

Problems Resolved in Release C.09.19

- **Bus Error (PR_83099)** — Bus error: HW Addr=0x29283030 IP=0x002086ac Task='mSnmprCtrl' Task ID=0x165ae00.
- **CDP/DISL (PR_92800)** — The switch does not forward CDP or DISL packets when CDP is disabled.
- **Counters (PR_87569)** — The ifHC ('High Capacity') MIB objects, for example ifHCInOctets, will return a maximum of 32 bits, rather than the correct 64 bits.
- **Crash (PR_87763)** — SNMP related memory leak where if the switch loses enough memory, it will crash with a message similar to:

SubSystem 0 went down: 04/17/02 08:07:45 Software exception at malloc.c:270 -- in 'mAdMgrCtrl', task ID = 0xa04ac0 -> Out of memory

or

SubSystem 0 went down: 04/04/02 14:23:42 Bus error: HW Addr=0x700a041c IP=0x00122ca0 PCW:0x00000003 ACW:0x00001001 Task='eDrvPoll' Task ID=0xaf17f0 pfp: 0x00af1cf0 sp:0x00af1d70 rip:0x00122ca0
- **IGMP (PR_87564)** — Adjusted interval between IGMP Group-Specific Queries (GSQs) in the normal Leave case when we are Querier.
- **IGMP (PR_87564)** — Switch improperly checks whether the value in the IP DA and/or Group Address is actually a legal IP Mcast before taking IGMP action on it.
- **IGMP (PR_87564)** — Updated the mechanism for what IP address we use when we send out Group-Specific Queries.
- **IGMP (PR_87564)** — When an IGMPv3 Join has one of the Reserved IP Mcast addresses in it the packet is double-forwarded, eventually causing the switch to run out of packet buffers.
- **IGMP (PR_87768)** — Enhancement to the IGMP Receive-Query code causing the switch to identify and ignore any Proxy Queries. Regular Queries will be treated the same as always.
- **IGMP (PR_87925)** — Enhancement for IGMPv3 support.
- **Meshing (PR_90509)** — Under certain conditions, loops can occur in a mesh topology.
- **Meshing (PR_91807)** — If a meshed link goes down, the meshed switches may not learn any new DAs until the switches are reset.
- **Meshing / Performance (PR_87904)** — When a mesh link is over-subscribed, meshing is migrating traffic too slowly from that link to an alternate link that is not over-subscribed.

- **SNMP (PR_88645)** — Added support for sending traps with non-PUBLIC community names.
- **SNMP/crash (PR_87568)** — An SNMP walk of the object cdpCacheDeviceId using an ifIndex value of 92 (Default VLAN) crashes the switch with a message similar to:

```
Bus error: HW Addr=0x3420504c IP=0x000d5e50 PCW:0x00000003  
ACW:0x00001004 Task='mSnmpCtrl' Task ID=0x84a9e0 pfp: 0x0084b010  
sp:0x0084b090 rip:0x000d5e50
```
- **SNTP (PR_87773)** — The SNTP Server remains at version 3 regardless of the configuration setting.
- **SNTP / Performance (PR_89815)** — Setting the SNTP broadcast mode from time sync to Unicast or TIMEP will result in CPU usage spiking to 99% and free packet buffers eventually depleting down to zero.
- **TELNET (PR_87599)** — If the user repeatedly portscans the management IP address of the switch, the user may no longer be able TELNET into the switch.
- **TimeP (PR_87602)** — The TimeP configuration is disabled after upgrading from C.08.XX to C.09.XX.
- **VTP (PR_87565)** — Changes made to not re-forward VTP packets if they arrive on a blocked port.
- **Web/Config (PR_87770)** — Web interface allows a user to create a VLAN with an ID of zero.

Release C.09.18

Problems Resolved in Release C.09.18

- **100/1000Base-T module (PR_92257)** — The HP J4115B 100/1000Base-T module may have connectivity problems with certain Intel Pro/1000 and other Broadcom 5703 or 5704 based NICs. These problems are often exhibited when the affected NIC workstation is either power cycled or the NIC driver is disabled then re-enabled. During the failure, the link LED may be lit, although no frames will be received by either side.

Release C.09.16

Problems Resolved in Release C.09.16

- **Address Learning/Forwarding Problems (Switch_4545)** — When monitoring is enabled on a switch and subsequently disabled, switch ports that were inactive while monitoring was enabled will not learn new MAC addresses after monitoring is disabled. Also, forwarding does not occur on these ports.

Release History

Release C. 09.14 (BETA)

- **Agent Hang (Switch_4523)** — The Agent hangs if 28 or more ASCII characters, (same or different ASCII characters), are entered into the password field at login access to the switch.
- **CDP (Switch_4511)** — Mesh ports drop incoming CDP packets.
- **CDP (Switch_4546)** — When Cisco Discovery Protocol (CDP) is disabled, the switch does not forward CDP packets.
- **Link-up polling interval (Switch_4543, 4548)** — A delay of up to 1 second between plugging in a cable (linkbeat established) and traffic being forwarded to and from that port may cause problems with some time sensitive applications. For example, AppleTalk dynamic address negotiation can be affected, resulting in multiple devices using the same AppleTalk address.
- **Meshing (Switch_4540)** — Balancing of traffic from an oversubscribed mesh link to an alternate mesh link is slow in comparison to earlier code releases.
- **Meshing (Switch_4568)** — Certain mesh configurations may cause switch mesh ports to not properly establish link with their mesh neighbors. This will prevent the switch from establishing itself into the mesh. The switch's event log will also improperly indicate that ldbal (mesh) ports are established.
- **Performance (Switch_4512)** — Slow performance and possible packet loss when switch is connected to Intel 10/100 NICs.
- **SNTP/TimeP (Switch_4559)** — Polling interval units not updated when changing configuration from SNTP (seconds) to TimeP (minutes) via the menu.
- **TFTP (Switch_4534)** — A TFTP PUT of a browse file larger than approximately 80K may starve other processes in the switch while the TFTP transfer is being processed. Symptoms include Telnet timeouts, Ping timeouts, and toggling of mesh ports during the TFTP transfer.
- **VTP (Switch_4518)** — Cisco's VTP (VLAN Trunk packets) are not forwarded.
- **WEB-Browser Interface (Switch_4519)** — IP Stack Member switches can be accessed without the use of a password.

Release C. 09.14 (BETA)

Problems Resolved in Release C.09.14

- **Agent Hang (Switch_4524)** — Agent processes (such as console, telnet, STP, ping, etc.) may stop functioning. This agent hang has been associated with the CERT SNMPv1 "encoding" test #1150.

(Various) Packet Buffer Leak: If the MAC Address Table is full, the switch may exhibit any of the following problems:

1. MAC addresses don't age out.

2. New MAC addresses are not learned.
3. Switch runs out of Packet Buffers.
4. Switch stops forwarding traffic.

Release C.09.13

Problems Resolved in Release C.09.13

- **Agent Hang (Switch_4515)** — Agent processes (such as console, telnet, STP, ping, etc.) may stop functioning when the IGMP querier function is disabled, and then re-enabled, on a VLAN that does not have an IP address configured.
- **Buffer Leak (Switch_4517)** — If IGMP is enabled, the switch may run out of packet buffers under certain conditions.

Release C.09.12 (BETA)

Problems Resolved in Release C.09.12

- **CDP (Switch_4511)** — The switch ignores CDP packets received on a mesh port. Therefore, the CLI command "walkmib cdpcacheentry" does not display mesh neighbor information.
- **IGMP (Switch_4516)** — Interoperability issues with some Cisco devices (such as some Cisco Catalyst 5000 & 6000 series switches) cause IGMP groups to be aged out of the switch's IGMP tables prematurely.
- **Stack overrun (Switch_4498)** — Code Red worm causes the switch to overrun its stack. [There is no external switch symptom to the Code Red worm.]
- **XPUT/PUT (SWITCH_4482)** — The CLI help text shows the wrong options for the XPUT and PUT commands.

Release C.09.10

Problems Resolved in Release C.09.10

- **Crash/Bus Error (SWITCH_4514)** — An SNMP Get request of a specific long OID can result in a bus error, an agent hang, or a switch crash with a message similar to:

```
-> Software_exception at svc_misc.s:379 -- in mCdpCtrl  
malloc_else_fatal() ran out of memory
```

This crash has been associated with the CERT SNMPv1 "Req-app" test suite.

Release C.09.09

Problems Resolved in Release C.09.09

- **Crash (SWITCH_4479)** — With meshing enabled, the switch may crash with a message similar to:

```
Software exception at memrpt.c:440 -- in 'mSnmpCtrl', task ID = 0x835f40
```

Release C.09.08 (BETA)

Problems Resolved in C.09.08

- The switch's CDP packets have been modified to better interoperate with older Cisco IOS versions. Certain legal CDP packets sent from the ProCurve switch could result in Cisco routers, running older IOS versions, to crash.

Note: The ProCurve switch's CDP packets are legal both before and after this modification.

- **Crash (SWITCH_4492)** — With meshing enabled, the switch may crash when a packet buffer is freed more than once, with a message similar to:

```
Software exception at alloc_free.c - in 'eDrvPoll' buf_free: corrupted  
buffer
```

- **Meshing (SWITCH_4493)** — MAC addresses that should be aged out are sometimes still in the address table when meshing is enabled. This does not result in connectivity problems.

Release C.09.07 (BETA)

- Minor modifications to internal switch functions.

Release C.09.06 (BETA)

Problems Resolved in Release C.09.06

- The sensitivity of the switch has been adjusted in detecting a device that exhibits an intermittent jabbering state. This results in the frequency of port online/offline toggling messages in the event log to be reduced.
- **IGMP (SWITCH_2997)** — The Switch stops sending IGMP Queries on some Vlans when Querier.

- **Telnet (SWITCH_2704)** — After the inactivity timer disconnects a TELNET session, subsequent TELNET sessions hang at the copyright screen.
- **Telnet/Console (SWITCH_2501)** — If the inactivity timer terminates a TELNET session, the switch will not allow another TELNET session. If the activity timer terminates a console session, the switch will not allow another console session.

Release C.09.05 (BETA)

Problems Resolved in Release C.09.05

- **Buffer Leak (SWITCH_2738)** — If CDP is enabled (by default), the switch runs out of packet buffers causing the switch to stop responding.
- **Buffer Leak (SWITCH_4471)** — A message buffer leak occurs when the switch receives a TACACS+ 'DISC' character.
- **CDP (SWITCH_4468)** — The switch sends the wrong MAC address for itself in CDP packets.
- **Memory Leak and Crash (SWITCH_4472)** — If the "Send Authentication Traps" trap receiver parameter on a Member switch is set to "Yes", it will cause a memory leak on the Stack Commander switch. The memory leak can eventually cause a crash. The specific details of the crash vary.
- **SNTP (SWITCH_4466)** — In the "Internet (IP) Service configuration screen", the value for the "SNTP Unicast Server Address" parameter is displayed incorrectly after a reboot.
- **SNTP (SWITCH_4467)** — The SNTP server IP address is listed as 0.0.0.0 in the "browse" output after a reboot. The value in the "config" output is correct.
- **Stack Management (SWITCH_4476)** — The Commander switch hangs when the switch is queried by the WhatsUp Gold utility.
- **Telnet (SWITCH_4458)** — The command "xget config" from a telnet session seems to "lock" the configuration file until reboot.
- **Web-browser Interface (SWITCH_2758)** — The Stack Management closeup button in the web browser does not display the Switch 4108GL.

Release C.09.04 (BETA)

Problems Resolved in Release C.09.04

- **SNTP (SWITCH_4465)** — SNTP class B or C unicast addresses are displayed as a very large value in the first octet.

Release C.09.03 (BETA)

Problems Resolved in Release C.09.03

- **CDP (SWITCH_2180)** — The System Information screen in the menu interface has no "Help" for CDP parameters.
- **Menu (SWITCH_4464)** — In the System Information configuration screen, the user cannot get the cursor to go to the "MAC Age Interval" field using the up or down arrow keys.
- **SNMP (SWITCH_4460)** — Modifying the MIB object dot1dStpPortPathCost on a port whose link is down causes the port to lock up and the link never to come up. Modifying this MIB object on a port which already has its link up does not cause link problems.
- **Stack Management (SWITCH_4463)** — Cannot access SNMP MIB objects of stack member switches using the SNMP v2c protocol.
- **TACACS+ (SWITCH_2169)** — In the Help screen for "TACACS+ Operation", the "Possible Values" for the Encryption Key is stated as "Up to ?? alphanumeric characters". The value "??" was replaced with 58, the actual number of characters allowed.

Release C.09.02

Problems Resolved in Release C.09.02

- Adds the following new features (documented in the "Software Update C.09.xx Release Notes", available on the WWW):
 - TACACS+
 - CDP (Cisco Discovery Protocol)
 - SNMP
 - Improved IGMP Capabilities
- Another new feature is CIDR Support. This feature allows the switch to be assigned an address in the "all zeros subnet" and in the "all ones subnet". Also supported is the ability to set the subnet mask to a supernet value (CIDR). This is a subnet mask that is shorter than the class default mask.

NOTE: The config file saved under version C.09.XX is NOT backward-compatible with previous software versions. The user is advised to save a copy of the pre-09.02 config file BEFORE UPGRADING to C.09.02 or greater, in case there is ever

a need to revert back to pre-09.02 software. Instructions for saving a copy of the config file are found in the "File Transfers" chapter in the Management and Configuration Guide that shipped with the switch.

- **Crash (SWITCH_4450)** — Switch crashes with a message similar to:

```
Software exception at ldbal_util.c:3165 -- in 'mIcmpCtrl'
```
- **Meshing (SWITCH_4457)** — Address aging of meshed switches is not synchronized.
- **Stack Management (SWITCH_4449,4451)** — Cannot add a candidate to a stack via the web-browser interface. Also, when the switch is a Commander a Member switch cannot be seen in the stack close-up view if that Member switch is a Switch 2512 or 2524.
- **STP (VARIOUS)** — Resolves several issues with STP, including: setting STP parameters via SNMP disables Spanning Tree, ability to configure the (reserved) spanning tree multicast address as a multicast filter address, bus error, and a switch crash with a message similar to:

```
-> Software exception at stp_ctrl.c:154 -- in 'mStpCtrl'
```

Release C.08.23 (BETA)

Problems Resolved in Release C.08.23

- **Auto-10 (SWITCH_4429)** — When a switch port configured for "auto-10" is connected to a NIC configured for 100 half duplex or 100 full duplex, the switch port comes up at 100half instead of 10half. [Fix is: switch port comes up at 10 Mbps. Result: switch detects a speed mismatch, disables the port, and generates an event log message similar to: fault: port 7 - Disabled due to speed mismatch. See help.]
- **Cisco RSM Interaction (SWITCH_4435)** — Due to a bug in Cisco's Routing Switch Module (RSM) software, when a ProCurve switch ARPs for the Cisco RSM's IP address, the Cisco learns the encapsulation type as "SNAP" - and then drops packets destined for the ProCurve switch. [Modification is: ProCurve switch's last-sent ARP is encapsulation type Ethernet II.] [NOTE: Cisco has new RSM software that fixes their bug - although the fix is not documented by Cisco; our information is that Cisco IOS version 12.1(4) for the RSM has the fix.]
- **Hotswap Crash (SWITCH_4434)** — When a module is removed from a switch slot, and a module with different port count is inserted into that slot, the switch may crash with an NMI message similar to:

```
NMI occurred: IP=0x00063c6c PCW:0x001c0003 ACW:0x00001002  
Task='mSnmpCtrl' Task ID=0x85bd40
```

- **TFTP (SWITCH_4432)** — Some error messages are not displayed in the event log.

Release C.08.22

Problems Resolved in Release C.08.22

- **Auto-TFTP (SWITCH_4423)** — Change made to streamline the Auto-TFTP operation.

Release C.08.21 (BETA)

- Update to the Manufacturing-only test routine.

Release C.08.20 (BETA)

- Adds new feature - support for TFTP upload of the software version specified in the switch's configuration file.

Release C.08.04 (BETA)

Problems Resolved in Release C.08.04

- **DHCP/Bootp (SWITCH_4415)** — When a switch is configured to use DHCP/Bootp to obtain an IP address, and the DHCP/Bootp server defines a config file for the switch to get via TFTP - if the config file has that same setting (use DHCP/Bootp to obtain an IP address) the switch will go into a loop, continually rebooting and getting the config file via TFTP.
- **ProCurve Stack Management(several issues) (SWITCH_4414)** — 1) SNMP stack access fails from outside the local network, 2) Discovery protocol is limited to 10 hops, 3) Possible memory corruption/memory leaks.
- **Web-browser interface (SWITCH_4416)** — Cannot select a port to monitor on slots E through J of the Switch 4000M/8000M.

Release C.08.03

Problems Resolved in Release C.08.03

- **Web-browser interface (SWITCH_4392)** — When using ProCurve Stack Management to communicate with a Stack Member, the web-browser interface appears to "hang" if connectivity is lost between Commander and Member.

Release C.08.02 (BETA)

Problems Resolved in Release C.08.02

- **Stack Management (SWITCH_4389)** — Default value for "auto-grab" was set to YES.
- **Stack Management (SWITCH_4390)** — Using SNMP to set the Stack Name on a Stack Member fails.

Release C.08.01 (BETA)

- Adds the following new features (documented in the "Software Update C.08.XX Release Notes", available on the WWW):
 - ProCurve Stack Management
 - VLAN configuration via web-browser interface
 - Auto-10 port configuration
 - Enhanced IP multicast filtering

NOTE: The config file saved under version C.08.01 is NOT backward-compatible with previous software versions. The user is advised to save a copy of the pre-08.01 config file BEFORE UPGRADING to C.08.01 or greater, in case there is ever a need to revert back to pre-08.01 software. Instructions for saving a copy of the config file are found in the "File Transfers" chapter in the Management and Configuration Guide that shipped with the switch.

Release C.07.28 (BETA)

Problems Resolved in C.07.28

- **Console (SWITCH_4312)** — Module Information screen shows "Module Type" of "1000T" instead of "100/1000T".
- **Flow Control (SWITCH_4304)** — Flow control cannot be configured on the 10Base-FL and 100Base-FX modules. (NOTE: This problem was introduced in C.07.20.)
- **XMODEM (SWITCH_4302)** — The binary "crashrec" file cannot be retrieved via XMODEM.

Release C.07.27

Problems Resolved in Release C.07.27

- Adds support for the J4115B ProCurve 100/1000Base-T (Gigabit-T) Module, model B.
- **ABC (SWITCH_4288)** — A "gratuitous ARP" is not forwarded when ABC is enabled for IP. This affects features such as MC ServiceGuard and Routing Switch redundant management module failover, because the new MAC-to-IP mapping must be forwarded to all endnodes that need to learn of the change.
- **Crash/SNMP (SWITCH_4224)** — A "walkmib" performed on the FEC trunk table can crash the switch. Also, the FEC trunk index may be wrong for a given FEC trunk number.

Release C.07.26

Problems Resolved in C.07.26

- **Switch "lockup" (SWITCH_3783)** — When a switch port configured for 10Mbps half-duplex operation is connected to a device configured for 100Mbps (half or full) operation, the switch port incorrectly shows link enabled and constant activity. After a short time interval in this condition (usually between 10 to 30 minutes) the switch stops forwarding traffic on all other (correctly-configured) ports. [Fix is to disable the port and generate an event log message similar to: fault: port 7 - Disabled due to speed mismatch. See help.]

Release C.07.25

Problems Resolved in Release C.07.25

- **Crash (SWITCH_4082)** — The switch may reboot when a user adds a new VLAN name via a telnet session.

Release C.07.24

Problems Resolved in Release C.07.24

- Adds new features: Support for 400 RMON Alarm Table traps, and support for 100 total traffic/security filters.
- **1000T Counters (SWITCH_3479)** — When 1000Base-T autonegotiates or loses connection, the counters are erroneously set to zero.
- **Event Log & Meshing (SWITCH_3902)** — After a mesh port recovers from a topology error (e.g. a hub connected between two switches configured for meshing), there is no event log entry to indicate that the mesh link is re-enabled.

- **IP Multicast (SWITCH_3804)** — Multicast streams do not stop for approximately 5 seconds after a "leave". [Fix changes that to 2 seconds, to comply with the spec.]
- **Meshing & Addresses (SWITCH_3870)** — The mesh will not learn a new MAC address if the incoming packet's source MAC address is equal to the destination MAC address. (Some UNIX systems send out such a packet upon bootup.)
- **Meshing & Broadcasts (SWITCH_3873)** — The first broadcast packet from any device is not correctly forwarded through a mesh. This can result in DHCP failures when the server is across a mesh from an endnode, for example.
- **Token Ring Cabling (SWITCH_3872)** — Spanning Tree Protocol (STP) does not correctly identify a "loopback" situation on an individual port. For example, if a twisted-pair port uses token ring cabling, and that cable is not connected to a live device, the token ring connector creates a short-circuit from transmit to receive. [Fix is that STP will identify this as a loop, and block the port.]

Release C.07.23

Problems Resolved in Release C.07.23

- **SNMP (SWITCH_3575)** — The MIB2 IP route table can become corrupted, causing problems when multiple VLANs are configured. (Related to SWITCH_3473.)

Release C.07.22

Problems Resolved in Release C.07.22

- **VLANs (SWITCH_3473)** — The configuration of 30 VLANs that use DHCP can cause the switch to become unable to forward traffic to the VLANs.

Release C.07.21 (BETA)

Problems Resolved in Release C.07.21

- **VLANs (SWITCH_3487,3495)** — Some dynamic changes to VLAN configuration are not properly displayed on the configuration screens.

Release C.07.20 (BETA)

Problems Resolved in Release C.07.20

- Adds support for the J4115A ProCurve 100/1000Base-T (Gigabit-T) Module. C.07.20 adds new features (documented in the updated "Management and Configuration Guide", available on the WWW), including:
 - Dynamic changes to any VLAN configuration
 - Support for ABC, switch meshing, FEC, and protocol-based filtering to the ProCurve Switches 2400M, 2424M, and 4000M
 - Support for 3 simultaneous telnet sessions
 - Timeout for telnet sessions
- **ABC (SWITCH_3453)** — Proxy ARP feature of Automatic Broadcast Control no longer functions (due to previous SWITCH_3410 fix).
- **Flow control (SWITCH_3418)** — Given a 10/100 port configured to auto-negotiate with flow control enabled, if that port auto-negotiates to half duplex, flow control is erroneously enabled.
- **XMODEM (SWITCH_3442)** — The Switch 2424M's console inactivity timer (with default value = 10 minutes) causes XMODEM download of switch software to fail. [NOTE: This bug fix description was erroneously omitted from the "readc723.txt" and earlier bug fix lists.]

Release C.07.02

Problems Resolved in Release C.07.02

- **ABC (SWITCH_3410)** — With Automatic Broadcast Control enabled for IP, unanswered ARPs generate extra broadcast traffic.
- **Command prompt (SWITCH_3413)** — Switch crashes if too many characters are entered at the command prompt.
- **IP (SWITCH_3424)** — Switch does not properly handle an IP "all subnets" broadcast, and may send ICMP "unreachable" messages.
- **SNMP (SWITCH_3411)** — The switch is incorrectly discovered as a router by OpenView Network Node Manager.
- **Web-browser interface (SWITCH_3404)** — Port configuration and port status are not updated, when the configuration changes are made via the web-browser interface.

- **Web-browser interface (SWITCH_3405)** — If a user with "Operator" password disables a port via the web-browser interface, the web-browser Configuration/Device View incorrectly shows that port as disabled. (Operator privileges do not allow that change, and the port is actually still enabled.)

Release C.07.01

Problems Resolved in Release C.07.01

- Adds the following new features (documented in the "Software Update C.07.XX Release Notes", available on the WWW):
 - Port Security
 - Authorized IP Managers
 - Class of Service (CoS)
 - New Command-Line Interface (CLI) commands (see "Read Me First", on the WWW)
- **10/100Base-T Module (SWITCH_3398)** — Newer versions of the J4111A ProCurve Switch 10/100Base-T Modules use a pair of 2MB memory chips, which are incorrectly interpreted as a pair of 1MB chips. The result is slow performance through that module.
- **ABC/VLAN (SWITCH_3193)** — ABC for IP does not function on a VLAN, if IP was configured after ABC on that VLAN.
- **Browse (SWITCH_2816)** — Cannot TFTP (using "put" command) the browse from the switch.
- **Command prompt (SWITCH_3310)** — Typing "help exit" causes the switch to exit back to the Diagnostics Menu.
- **Config (SWITCH_3354)** — When a switch module is removed and the switch is not rebooted, port configuration changes cannot be saved.
- **Console (SWITCH_3284)** — When the console session is disconnected by the configured "timeout" period, there is no message to the user.
- **Crash (SWITCH_3335)** — Switch crashes with this message: system: Fault:03 subType:02 IP:0x00189358 PCW:0x00000003The following combination of events can cause this crash:
 1. Switch module is replaced with a different module and switch is not rebooted
 2. Switch is accessed via the web-browser interface
- **Filters (SWITCH_3277)** — Protocol filter option "NetBIOS" should be "NetBEUI".
- **IGMP (SWITCH_3218)** — Switch does not identify PIM routers that comply with the new PIM specification which no longer uses IGMP packets.

Release History

Release C.06.07 (BETA)

- **IGMP (SWITCH_3351)** — A "join" packet sent to the reserved range of 01-00-5e-00-00-xx is dropped, causing problems with protocols like Cisco's HSRP (which use that reserved range, whether or not they should).
- **Meshing (SWITCH_3149)** — After configuring a mesh port, the IGMP Service and VLAN Port Assignment configuration screens do not list the newly-configured mesh port until switch is rebooted.
- **SNMP (SWITCH_1415)** — After configuring a trap receiver a warm start trap is sent out, even though the switch has not rebooted.
- **VLAN (SWITCH_3023)** — Outbound tagged packets are always sent with priority = 0.
- **Walkmib (SWITCH_2935)** — If more than one filter is configured, walkmib does not always display all the filter parameters.

Release C.06.07 (BETA)

Problems Resolved in Release C.06.07

- **Gigabit Stacking Module (SWITCH_3309)** — If the transceiver type is changed in the module, the transceiver may not re-establish link after the module is hotswapped (for the J4093A ProCurve Switch 2424M only).
- **IGMP (SWITCH_3361)** — IGMP "join" reports are not forwarded to routers if IGMP is disabled and then re-enabled.

Release C.06.06

- Modification of Lab troubleshooting commands.
(These commands are not available to the switch operator.)

Release C.06.05 (BETA)

Problems Resolved in Release C.06.05

- **Bus error (SWITCH_3264)** — The switch may crash with a bus error if it receives an IGMP packet with an invalid IP header length.
- **Telnet (SWITCH_3278)** — When viewing the switch's event log from a telnet session, the session may halt ("lock up").

Release C.06.04 (BETA)

Problems Resolved in C.06.04

- **Broadcast Limit, ABC, and Hot-swap (SWITCH_3266,3272)** — Given this situation:

1. a new or different module is hot-swapped into a switch slot, and
2. the switch is rebooted.

These problems are seen:

- a. if ABC is enabled, broadcast limit is set to 0 instead of 30, and
 - b. if ABC is disabled, broadcast limit is set to 30 instead of 0
- **Crash (SWITCH_3267)** — The switch may crash with an NMI error when a J4130A Switch 2424M Gigabit Stacking Module is inserted into a J4093A ProCurve Switch 2424M that is powered on.

Release C.06.03

- Adds support for the J4116A Switch 2424M Gigabit Stacking Kit (for the J4093A ProCurve Switch 2424M only).

Release C.06.02 (BETA)

Problems Resolved in Release C.06.02

- Adds support for the J4130A Switch 2424M Gigabit Stacking module, the J4131A ProCurve Gigabit-SX transceiver, and the J4132A ProCurve Gigabit-LX transceiver (for the J4093A ProCurve Switch 2424M only).

- **Console (SWITCH_3249)** — If the console cable is connected to a terminal or PC for long periods of time, the switch may crash with the following message:

```
Software exception in ISR at sysLib.c:498
```

```
-> Reboot through sysToMonitor
```

- **IP connectivity (SWITCH_3239)** — Multiple failed attempts to contact a timep server can deplete system buffers ("mbufs") and cause IP connectivity to fail to/from the switch.
- **LEDs (SWITCH_3252)** — Power supply LEDs are reversed. The LED for power supply #2 lights when power supply #1 is installed, and vice versa.

Release History

Release C.06.01

- **Selftest (SWITCH_3247)** — Power-cycle selftest does not properly indicate a Switch Engine failure. However, the failure is properly indicated when rebooting the switch using the "reset" button.
- **Timep (SWITCH_3242)** — If a negative number is configured in the "Time Zone" field (to denote minutes East of GMT), two things happen:
 1. the negative number is ignored, giving a switch time equal to GMT, and
 2. the switch does not poll the Timep server after the initial time setting.
- **Timep and VLANs (SWITCH_3038)** — The switch can only reach a timep server that is on the first VLAN in the "VLAN Names" list. This is true even if that VLAN has been renamed from "DEFAULT_VLAN", or has a VLAN ID other than 1.
- **Timep enhancements (SWITCH_3112,3113)** — 1) allow dynamic reconfiguration of the timep server's IP address, and 2) add retries when the timep server does not respond.
- **Web-browser interface (SWITCH_3238)** — 10Base-FL module does not display in the "Closeup View" of a Switch 8000M (slot appears empty).

Release C.06.01

Problems Resolved in Release C.06.01

- First release for the J4093A ProCurve Switch 2424M.
- Also, adds new feature: Support for 30 VLANs on a switch. NOTE: The config file saved under version C.06.01 is NOT backward-compatible with previous software versions. The user is advised to save a copy of the pre-06.01 config file BEFORE UPGRADING to C.06.01 or greater, in case there is ever a need to revert back to pre-06.01 software. Instructions for saving a copy of the config file are found in the "File Transfers" chapter in the Management and Configuration Guide that shipped with the switch.
- **BOOTP (SWITCH_3179)** — The switch will go into a continual reboot loop, when ALL the following conditions are true:
 1. Switch is configured to use BOOTP, and
 2. BOOTP server is set up to download a config file to the switch, and
 3. The config file on the BOOTP server is configured to use BOOTP, and
 4. The software version running on the switch is different than the software version that created the config file on the BOOTP server.
- **Crash (SWITCH_3200)** — Switch crashes with this message:

```
system: Fault:02 subType:02 IP:0x00021be8 PCW:0x00000003
```

This crash has been associated with network management applications such as TopTools.

- **Crash (SWITCH_3214)** — Switch may crash with a "bus error" when hot-swapping a Gigabit-SX card.
- **Filters (SWITCH_3225)** — Under heavy traffic loads, dynamically modifying filters can cause those filters to malfunction.
- **IGMP (SWITCH_3199)** — In situations of very high traffic on a port, IGMP protocol packets (e.g. queries) are not sent by the switch.
- **Monitoring Port (SWITCH_3224)** — Dynamically adding/deleting/changing the "network monitoring port" causes trunks and filters to malfunction.
- **Selftest (SWITCH_3208)** — With a 10/100 EISA NIC in an HP-UX D-class server - if the server is plugged into the switch when the switch reboots, the switch module fails selftest.
- **VLANs (SWITCH_3235)** — Given several configured VLANs, the act of deleting a VLAN that is in the middle of the list (i.e. not the first or the last one), causes filters, network monitoring port, and meshing to malfunction.

Release C.05.07

Problems Resolved in Release C.05.07

- Adds new features: Spanning Tree "fast" mode, and the ability to enable the switch as a TFTP server to retrieve its config, using one of these case-sensitive commands:

```
setmib hpSwitchIpTftpMode.0 -i 1 (disable tftp server, the default)
```

```
setmib hpSwitchIpTftpMode.0 -i 2 (enable tftp server)
```

- **BOOTP (SWITCH_3134)** — Switch does not read the "T144" field in the bootptab file, so the switch fails to download its config file via BOOTP.
- **Crash (SWITCH_3181)** — Switch crashes with this message:

```
system: Fault:02 subType:02 IP:0x00000000 PCW:0x00002002
```

This crash has been associated with network management applications such as TopTools.

- **Web-browser interface (SWITCH_3216)** — Web-browser gives error "404 page not found" when applying changes in the "Configuration/Device Features" tab on a Switch 4000M that has no VLANs configured.

NOTE: Additional versions of C.05.xx were distributed during C.06.xx development. Therefore, bugs fixed in the first few versions of C.06.xx may have also been fixed in the additional C.05.xx versions. These bugs were fixed in the additional C.05.xx versions:

C.05.08	C.05.09	C.05.10
SWITCH_3179	SWITCH_3214	SWITCH_3238
SWITCH_3199	SWITCH_3224	SWITCH_3242
SWITCH_3200	SWITCH_3225	SWITCH_3247
SWITCH_3208	SWITCH_3235	

Also, C.05.10 includes C.05.09 fixes, and C.05.09 includes C.05.08 fixes.

Release C.05.06

Problems Resolved in Release C.05.06

- Adds support for the HP ProCurve Switch Gigabit-LX module (J4114A).
- **Crash (SWITCH_3168)** — Switch crashes with this message:

```
system: Fault: 03 subType:02 IP:0x000a950c PCW:0x00000003  
ACW:0x00001004
```
- **Crash (SWITCH_3169)** — Switch crashes with a message similar to:

```
NMI occurred: IP=0x000a9550 PCW:0x00000003 ACW:0x00001004
```

The "IP" values may be in the range 0x000a9550 - 0x000a956c.
- **Flow control (SWITCH_3151)** — With flow control enabled, switch drops multicast and broadcast packets.

Release C.05.05

Problems Resolved in Release C.05.05

- Resolves the problems listed below.
- **IP (SWITCH_3090)** — Switch does not reply to an IP address that it thinks is a broadcast destination (because switch wrongly uses its own mask to decide that the destination address is a broadcast).

- **Selftest (SWITCH_3110)** — Error codes explaining selftest failures are incorrect.
- **Web-browser Interface (SWITCH_3130)** — "Apply Changes" fails on "Configuration/Device Features" screen when there are no VLANs configured on the switch.

Release C.05.04

Problems Resolved in Release C.05.04

- Adds support for the HP ProCurve Switch 4000M/2400M.
- **Meshing (VARIOUS)** — Resolves several issues with switch meshing, involving memory, path cost, state (established or not), failed connectivity across the mesh, rebooting, and erroneous report of loops during extremely high traffic.
- **SNMP (SWITCH_3082)** — Switch does not give SNMP replies for PDUs greater than 484 bytes.

Release C.05.03

Problems Resolved in Release C.05.03

- Resolves the problems listed below.
- **Web-browser Interface (SWITCH_3160)** — Gigabit SX card in a Switch 1600M is not displayed.

Release C.05.02

- First release for the HP ProCurve Switch 8000M/1600M.



© 2001, 2004 Hewlett-Packard Development Company, LP. The information contained herein is subject to change without notice.

Edition 4, May 2004
Manual Part Number
5969-2375
