



Multicast and Routing Guide

2900

ProCurve Switches
T.11.XX

www.procurve.com



ProCurve Switch 2900

August 2006

T.11.xx

Multicast and Routing Guide

© Copyright 2006 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. All Rights Reserved.

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

5991-6199
August 2006

Applicable Products

ProCurve Switch 2900-24G (J9049A)
ProCurve Switch 2900-48G (J9050A)

Trademark Credits

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

Disclaimer

The information contained in this document is subject to change without notice.

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

The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Hewlett-Packard assumes no responsibility for the use or reliability of its software on equipment that is not furnished by Hewlett-Packard.

Warranty

See the Customer Support/Warranty booklet included with the product.

A copy of the specific warranty terms applicable to your Hewlett-Packard products and replacement parts can be obtained from your HP Sales and Service Office or authorized dealer.

Contents

Product Documentation

About Your Switch Manual Set	vii
Feature Index	viii

1 Getting Started

Contents	1-1
Introduction	1-2
Conventions	1-2
Feature Descriptions by Model	1-2
Command Syntax Statements	1-3
Command Prompts	1-3
Screen Simulations	1-4
Port Identity Examples	1-4
Configuration and Operation Examples	1-4
Keys	1-4
Sources for More Information	1-5
Getting Documentation From the Web	1-7
Online Help	1-7
Need Only a Quick Start?	1-8
IP Addressing	1-8
To Set Up and Install the Switch in Your Network	1-9
Physical Installation	1-9

2 Multimedia Traffic Control with IP Multicast (IGMP)

Contents	2-1
Overview	2-2
IGMP General Operation and Features	2-3

IGMP Terms	2-4
IGMP Operating Features	2-5
Basic Operation	2-5
Enhancements	2-5
Number of IP Multicast Addresses Allowed	2-6
CLI: Configuring and Displaying IGMP	2-7
How IGMP Operates	2-12
Operation With or Without IP Addressing	2-13
Automatic Fast-Leave IGMP	2-14
Forced Fast-Leave IGMP	2-17
Configuring Delayed Group Flush	2-18
IGMP Proxy Forwarding	2-18
How IGMP Proxy Forwarding Works	2-18
CLI Commands for IGMP Proxy Configuration	2-20
VLAN Context Command	2-21
IGMP Proxy Show Command	2-22
Operating Notes for IGMP Proxy Forwarding	2-23
Using the Switch as Querier	2-26
Excluding Well-Known or Reserved	
Multicast Addresses from IP Multicast Filtering	2-27

3 IP Routing Features

Contents	3-1
Overview of IP Routing	3-3
IP Interfaces	3-4
IP Tables and Caches	3-4
ARP Cache Table	3-5
IP Route Table	3-5
IP Forwarding Cache	3-6
IP Global Parameters for Routing Switches	3-7
IP Interface Parameters for Routing Switches	3-9
Configuring IP Parameters for Routing Switches	3-10
Configuring IP Addresses	3-10
Configuring ARP Parameters	3-10

How ARP Works	3-10
Enabling Proxy ARP	3-12
Configuring Forwarding Parameters	3-13
Changing the TTL Threshold	3-13
Enabling Forwarding of Directed Broadcasts	3-13
Configuring ICMP	3-14
Disabling ICMP Messages	3-14
Disabling Replies to Broadcast Ping Requests	3-14
Disabling ICMP Destination Unreachable Messages	3-15
Disabling ICMP Redirects	3-16
Configuring Static IP Routes	3-16
Static Route Types	3-16
Other Sources of Routes in the Routing Table	3-17
Static IP Route Parameters	3-17
Static Route States Follow VLAN States	3-18
Configuring a Static IP Route	3-18
Displaying Static Route Information	3-20
Configuring the Default Route	3-20
Configuring IRDP	3-21
Enabling IRDP Globally	3-22
Enabling IRDP on an Individual VLAN Interface	3-22
Displaying IRDP Information	3-23
Configuring DHCP Relay	3-24
Overview	3-24
DHCP Option 82	3-24
Introduction	3-24
Option 82 Server Support	3-26
Terminology	3-26
General DHCP Option 82 Requirements and Operation	3-27
Option 82 Field Content	3-28
Forwarding Policies	3-30
Multiple Option 82 Relay Agents in a Client Request Path	3-32
Validation of Server Response Packets	3-33
Multinetted VLANs	3-34
Configuring Option 82 Operation on the Routing Switch	3-36

Operating Notes	3-37
DHCP Packet Forwarding	3-38
Unicast Forwarding	3-38
Broadcast Forwarding	3-38
Minimum Requirements for DHCP Relay Operation	3-39
Enabling DHCP Relay	3-39
Configuring a Helper Address	3-39
Viewing the Current DHCP Relay Configuration	3-40
UDP Broadcast Forwarding	3-41
Overview	3-41
Subnet Masking for UDP Forwarding Addresses	3-42
Configuring and Enabling UDP Broadcast Forwarding	3-43
Globally Enabling UDP Broadcast Forwarding	3-43
Configuring UDP Broadcast Forwarding on Individual VLANs .	3-43
Displaying the Current IP Forward-Protocol Configuration	3-45
Operating Notes for UDP Broadcast Forwarding	3-46
Messages Related to UDP Broadcast Forwarding	3-46
Index	1

Product Documentation

About Your Switch Manual Set

The switch manual set includes the following documentation:

- *Read Me First*—a printed guide shipped with your switch. Provides software update information, product notes, and other information.
- *Installation and Getting Started Guide*—a printed guide shipped with your switch. This guide explains how to prepare for and perform the physical installation and connect the switch to your network.
- *Management and Configuration Guide*—a PDF on the ProCurve Networking Web Site that describes how to configure, manage, and monitor basic switch operation.
- *Advanced Traffic Management Guide*—a PDF on the ProCurve Networking Web Site that explains how to configure traffic management features such as VLANs, MSTP, and QoS.
- *Multicast and Routing Guide*—a PDF on the ProCurve Networking Web Site that explains how to configure IGMP and IP routing.
- *Access Security Guide*—a PDF on the ProCurve Networking Web Site that explains how to configure access security features and user authentication on the switch.
- *Release Notes*—posted on the ProCurve Networking Web Site to provide information on software updates. The release notes describe new features, fixes, and enhancements that become available between revisions of the main product guide.

Note

For the latest version of all ProCurve switch documentation, including Release Notes covering recently added features, visit the ProCurve Networking Web Site at www.procurve.com, click on **Technical support**, and then click on **Product manuals (all)**.

Feature Index

For the manual set supporting your switch model, the following feature index indicates which manual to consult for information on a given software feature.

Feature	Management and Configuration	Advanced Traffic Management	Multicast and Routing	Access Security Guide
802.1Q VLAN Tagging		X		
802.1p Priority	X			
802.1X Port-Based Authentication				X
AAA Authentication				X
Authorized IP Managers				X
Authorized Manager List (web, telnet, TFTP)				X
Auto MDIX Configuration	X			
BOOTP	X			
Config File	X			
Console Access	X			
Copy Command	X			
CoS (Class of Service)		X		
Debug	X			
DHCP Configuration		X		
DHCP Option 82			X	
DHCP/Bootp Operation	X			
Diagnostic Tools	X			
Downloading Software	X			
Eavesdrop Protection				X
Event Log	X			
Factory Default Settings	X			

Feature	Management and Configuration	Advanced Traffic Management	Multicast and Routing	Access Security Guide
Flow Control (802.3x)	X			
File Management	X			
File Transfers	X			
Friendly Port Names	X			
GVRP		X		
Identity-Driven Management (IDM)		X		
IGMP			X	
Interface Access (Telnet, Console/Serial, Web)	X			
IP Addressing	X			
IP Routing			X	
Jumbos Support		X		
LACP	X			
Link	X			
LLDP	X			
LLDP-Med	X			
MAC Address Management	X			
MAC Lockdown				X
MAC Lockout				X
MAC-based Authentication				X
MAC authentication RADIUS support				X
Management VLAN		X		
Monitoring and Analysis	X			
Multicast Filtering				X
Multiple Configuration Files	X			
Network Management Applications (SNMP)	X			
OpenView Device Management	X			
Passwords and Password Clear Protection				X

Product Documentation

Feature	Management and Configuration	Advanced Traffic Management	Multicast and Routing	Access Security Guide
PCM	X			
Ping	X			
Port Configuration	X			
Port Monitoring		X		
Port Security				X
Port Status	X			
Port Trunking (LACP)	X			
Port-Based Access Control				X
Port-Based Priority (802.1Q)	X			
Protocol Filters				X
Protocol VLANS		X		
Quality of Service (QoS)		X		
RADIUS Authentication and Accounting				X
RADIUS-Based Configuration		X		
RMON 1,2,3,9	X			
Routing			X	
Routing - IP Static			X	
Secure Copy	X			
SFLOW	X			
SFTP	X			
SNMPv3	X			
Software Downloads (SCP/SFTP, TFPT, Xmodem)	X			
Source-Port Filters				X
Spanning Tree (MSTP)		X		
SSHv2 (Secure Shell) Encryption				X
SSLv3 (Secure Socket Layer)				X
Stack Management		X		

Feature	Management and Configuration	Advanced Traffic Management	Multicast and Routing	Access Security Guide
Syslog	X			
System Information	X			
TACACS+ Authentication				X
Telnet Access	X			
TFTP	X			
Time Protocols (TimeP, SNTP)	X			
Traffic/Security Filters				X
Troubleshooting	X			
VLANs		X		
VLAN Mirroring (1 static VLAN)		X		
Web Authentication RADIUS Support				X
Web-based Authentication				X
Web UI	X			
Xmodem	X			

—This page is intentionally unused—

Getting Started

Contents

Introduction	1-2
Conventions	1-2
Feature Descriptions by Model	1-2
Command Syntax Statements	1-3
Command Prompts	1-3
Screen Simulations	1-4
Port Identity Examples	1-4
Configuration and Operation Examples	1-4
Keys	1-4
Sources for More Information	1-5
Getting Documentation From the Web	1-7
Online Help	1-7
Need Only a Quick Start?	1-8
IP Addressing	1-8
To Set Up and Install the Switch in Your Network	1-9
Physical Installation	1-9

Introduction

This *Management and Configuration Guide* is intended for use with the following switches:

- ProCurve Switch 2900-24G
- ProCurve Switch 2900-48G

This guide describes how to use the command line interface (CLI), Menu interface, and web browser to configure, manage, monitor, and troubleshoot switch operation.

For an overview of other product documentation for the above switches, refer to “*Product Documentation*” on page vii.

You can download documentation from the ProCurve Networking Web Site, www.procurve.com.

Caution

Use only the supported genuine ProCurve mini-GBICs with your switch. Non-ProCurve mini-GBICs are not supported.

Conventions

This guide uses the following conventions for command syntax and displayed information.

Feature Descriptions by Model

In cases where a software feature is not available in all of the switch models covered by this guide, the section heading specifically indicates which product or product series offer the feature.

For example, (the switch is highlighted here in ***bold italics***):

“QoS Pass-Through Mode on the ***Switch 2900***”.

Command Syntax Statements

Syntax: ip default-gateway < ip-addr >

Syntax: show interfaces [port-list]

- Vertical bars (|) separate alternative, mutually exclusive elements.
- Square brackets ([]) indicate optional elements.
- Braces (< >) enclose required elements.
- Braces within square brackets ([< >]) indicate a required element within an optional choice.
- Boldface indicates use of a CLI command, part of a CLI command syntax, or other displayed element in general text. For example:
 “Use the **copy tftp** command to download the key from a TFTP server.”
- Italics indicate variables for which you must supply a value when executing the command. For example, in this command syntax, you must provide one or more port numbers:

Syntax: aaa port-access authenticator < port-list >

Command Prompts

In the default configuration, your switch displays a CLI prompt similar to the following:

```
ProCurve 2900-24G#
```

To simplify recognition, this guide uses **ProCurve** to represent command prompts for all models. For example:

```
ProCurve#
```

(You can use the **hostname** command to change the text in the CLI prompt.)

Screen Simulations

Displayed Text. Figures containing simulated screen text and command output look like this:

```
ProCurve> show version
Image stamp:   /sw/code/build/info
               March 1, 2006 13:43:13
               T.11.01
               139

ProCurve>
```

Figure 1-1. Example of a Figure Showing a Simulated Screen

In some cases, brief command-output sequences appear without figure identification. For example:

```
ProCurve(config)# clear public-key
ProCurve(config)# show ip client-public-key
show_client_public_key: cannot stat keyfile
```

Port Identity Examples

This guide describes software applicable to both chassis-based and stackable ProCurve switches. Where port identities are needed in an example, this guide uses the chassis-based port identity system, such as “A1”, “B3-B5”, “C7”, etc. However, unless otherwise noted, such examples apply equally to the stackable switches, which typically use only numbers, such as “1”, “3-5”, “15”, etc. for port identities.

Configuration and Operation Examples

Unless otherwise noted, examples using a particular switch model apply to all switch models covered by this guide.

Keys

Simulations of actual keys use a bold, sans-serif typeface with square brackets. For example, the Tab key appears as **[Tab]** and the “Y” key appears as **[Y]**.

Sources for More Information

For additional information about switch operation and features not covered in this guide, consult the following sources:

- Feature Index—For information on which product manual to consult for a given software feature, refer to the “Feature Index” on page viii.

Note

For the latest version of all ProCurve switch documentation, including Release Notes covering recently added features, visit the ProCurve Networking Web Site at www.procurve.com, click on **Technical support**, and then click on **Product Manuals (all)**.

- Software Release Notes—Release notes are posted on the ProCurve Networking web site and provide information on new software updates:
 - new features and how to configure and use them
 - software management, including downloading software to the switch
 - software fixes addressed in current and previous releases

To view and download a copy of the latest software release notes for your switch, refer to “Getting Documentation From the Web” on page 1-7.

- Product Notes and Software Update Information—The printed *Read Me First* shipped with your switch provides software update information, product notes, and other information. For the latest version, refer to “Getting Documentation From the Web” on page 1-7.
- *Installation and Getting Started Guide*—Use the *Installation and Getting Started Guide* shipped with your switch to prepare for and perform the physical installation. This guide also steps you through connecting the switch to your network and assigning IP addressing, as well as describing the LED indications for correct operation and trouble analysis. You can download a copy from the ProCurve Networking web site. (See “Getting Documentation From the Web” on page 1-7.)


- *Management and Configuration Guide*—Use this guide for information on topics such as:
 - various interfaces available on the switch
 - memory and configuration operation
 - interface access
 - IP addressing
 - time protocols
 - port configuration, trunking, and traffic control
 - SNMP, LLDP, and other network management topics
 - file transfers, switch monitoring, troubleshooting, and MAC address management
- *Advanced Traffic Management Guide*—Use this guide for information on topics such as:
 - VLANs: Static port-based and protocol VLANs, and dynamic GVRP VLANs
 - Spanning-Tree: 802.1s (MSTP)
 - Quality-of-Service (QoS)
- *Multicast and Routing Guide*—Use this guide for information topics such as:
 - IGMP
 - IP routing
- *Access Security Guide*—Use this guide for information on topics such as:
 - Local username and password security
 - Web-Based and MAC-based authentication
 - RADIUS and TACACS+ authentication
 - SSH (Secure Shell) and SSL (Secure Socket Layer) operation
 - 802.1X access control
 - Port security operation with MAC-based control
 - Authorized IP Manager security
 - Key Management System (KMS)

Getting Documentation From the Web


1. Go to the ProCurve Networking Web Site at
www.procurve.com
2. Click on **Technical support**.
3. Click on **Product manuals**.
4. Click on the product for which you want to view or download a manual.

Online Help

If you need information on specific parameters in the menu interface, refer to the online help provided in the interface. For example:



If you need information on a specific command in the CLI, type the command name followed by “help”. For example:



If you need information on specific features in the ProCurve Web Browser Interface (hereafter referred to as the “web browser interface”), use the online help available for the web browser interface. For more information on web browser Help options, refer to “Online Help for the ProCurve Web Browser Interface” in the Management and Configuration Guide.

If you need further information on ProCurve switch technology, visit the ProCurve Networking web site at:

www.procurve.com

Need Only a Quick Start?

IP Addressing

If you just want to give the switch an IP address so that it can communicate on your network, or if you are not using VLANs, ProCurve recommends that you use the Switch Setup screen to quickly configure IP addressing. To do so, do one of the following:

- Enter **setup** at the CLI Manager level prompt.
`Procurve# setup`
- In the Main Menu of the Menu interface, select

8. Run Setup

For more on using the Switch Setup screen, see the *Installation and Getting Started Guide* you received with the switch.

To Set Up and Install the Switch in Your Network

Physical Installation

Use the ProCurve *Installation and Getting Started Guide* (shipped with the switch) for the following:

- Notes, cautions, and warnings related to installing and using the switch and its related modules
- Instructions for physically installing the switch in your network
- Quickly assigning an IP address and subnet mask, set a Manager password, and (optionally) configure other basic features.
- Interpreting LED behavior.

For the latest version of the *Installation and Getting Started Guide* for your switch, refer to “Getting Documentation From the Web” on page 1-7.

Getting Started

To Set Up and Install the Switch in Your Network

—This page is intentionally unused—

Multimedia Traffic Control with IP Multicast (IGMP)

Contents

Overview	2-2
IGMP General Operation and Features	2-3
IGMP Terms	2-4
IGMP Operating Features	2-5
Basic Operation	2-5
Enhancements	2-5
Number of IP Multicast Addresses Allowed	2-6
CLI: Configuring and Displaying IGMP	2-7
How IGMP Operates	2-12
Operation With or Without IP Addressing	2-13
Automatic Fast-Leave IGMP	2-14
Forced Fast-Leave IGMP	2-17
Configuring Delayed Group Flush	2-18
IGMP Proxy Forwarding	2-18
How IGMP Proxy Forwarding Works	2-18
CLI Commands for IGMP Proxy Configuration	2-20
VLAN Context Command	2-22
IGMP Proxy Show Command	2-23
Operating Notes for IGMP Proxy Forwarding	2-24
Using the Switch as Querier	2-26
Excluding Well-Known or Reserved Multicast Addresses from IP Multicast Filtering	2-27

Overview

This chapter describes multimedia traffic control with IP multicast (IGMP) to reduce unnecessary bandwidth usage on a per-port basis, and how to configure it with the switch's built-in interfaces:

For general information on how to use the switch's built-in interfaces, refer to these chapters in the *Management and Configuration Guide* for your switch:

- Chapter 3, "Using the Menu Interface"
- Chapter 4, "Using the Command Line Interface (CLI)"
- Chapter 5, "Using the ProCurve Web Browser Interface"
- Chapter 6, "Switch Memory and Configuration"

Note

The use of static multicast filters is described in the chapter titled "Traffic/Security Filters" in the *Access Security Guide* for your ProCurve switch.

IGMP General Operation and Features

IGMP Features

Feature	Default	Menu	CLI
view igmp configuration	n/a	—	page 2-7
show igmp status for multicast groups used by the selected VLAN	n/a	—	Yes
enabling or disabling IGMP (Requires VLAN ID Context)	disabled	—	page 2-9
per-port packet control	auto	—	page 2-10
IGMP traffic priority	normal	—	page 2-11
querier	enabled	—	page 2-11
fast-leave	disabled	—	page 2-14

In a network where IP multicast traffic is transmitted for various multimedia applications, you can use the switch to reduce unnecessary bandwidth usage on a per-port basis by configuring IGMP (Internet Group Management Protocol controls). In the factory default state (IGMP disabled), 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. Enabling IGMP allows the ports to detect IGMP queries and report packets and manage IP multicast traffic through the switch.

IGMP is useful in multimedia applications such as LAN TV, desktop conferencing, and collaborative computing, where there is multipoint communication; that is, communication from one to many hosts, or communication originating from many hosts and destined for many other hosts. In such multipoint applications, IGMP will be configured on the hosts, and multicast traffic will be generated by one or more servers (inside or outside of the local network). Switches in the network (that support IGMP) can then be configured to direct the multicast traffic to only the ports where needed. If multiple VLANs are configured, you can configure IGMP on a per-VLAN basis.

Enabling IGMP allows detection of IGMP queries and report packets in order to manage IP multicast traffic through the switch. If no other querier is detected, the switch will then also function as the querier. (If you need to disable the querier feature, you can do so through the IGMP configuration MIB. Refer to “Changing the Querier Configuration Setting” on page 2-11.)

Note

IGMP configuration on the switches covered in this guide operates at the VLAN context level. If you are not using VLANs, then configure IGMP in VLAN 1 (the default VLAN) context.

IGMP Terms

- **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.
- **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. A querier uses data received 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. When enabled (the default state), the switch's querier function eliminates the need for a multicast router. In most cases, ProCurve recommends that you leave this parameter in the default "enabled" state even if you have a multicast router performing the querier function in your multicast group. For more information, see "How IGMP Operates" on page 2-12.

IGMP Operating Features

Basic Operation

In the factory default configuration, IGMP is disabled. To enable IGMP

- If multiple VLANs are not configured, you configure IGMP on the default VLAN (DEFAULT_VLAN; VID = 1).
- If multiple VLANs are configured, you configure IGMP on a per-VLAN basis for every VLAN where this feature is to be used.

Enhancements

With the CLI, you can configure these additional options:

- **Forward with High Priority.** Disabling this parameter (the default) causes the switch or VLAN to process IP multicast traffic, along with other traffic, in the order received (usually, normal priority). Enabling this parameter causes the switch or VLAN to give a higher priority to IP multicast traffic than to other traffic.
- **Auto/Blocked/Forward:** You can use the console to configure individual ports to any of the following states:
 - **Auto** (the default): Causes the switch to interpret IGMP packets and to filter IP multicast traffic based on the IGMP packet information for ports belonging to a multicast group. This means that IGMP traffic will be forwarded on a specific port only if an IGMP host or multicast router is connected to the port.
 - **Blocked:** Causes the switch to drop all IGMP transmissions received from a specific port and to block all outgoing IP Multicast packets for that port. This has the effect of preventing IGMP traffic from moving through specific ports.
 - **Forward:** Causes the switch to forward all IGMP and IP multicast transmissions through the port.
- **Operation With or Without IP Addressing:** This feature helps to conserve IP addresses by enabling IGMP to run on VLANs that do not have an IP address. See “Operation With or Without IP Addressing” on page 2-13.
- **Querier Capability:** The switch performs this function for IGMP on VLANs having an IP address when there is no other device in the VLAN acting as querier. See “Using the Switch as Querier” on page 2-26.

Notes

Whenever IGMP is enabled, the switch generates an Event Log message indicating whether querier functionality is enabled.

IP multicast traffic groups are identified by IP addresses in the range of 224.0.0.0 to 239.255.255.255. Also, 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 “Excluding Well-Known or Reserved Multicast Addresses from IP Multicast Filtering” on page 2-27.

For more information, refer to “How IGMP Operates” on page 2-12.

Number of IP Multicast Addresses Allowed

The total of IGMP filters (addresses) and static multicast filters together is 2047 if data driven or 2048 otherwise, depending on the current max-vlans configuration. If multiple VLANs are configured, then each filter is counted once per VLAN in which it is used.

CLI: Configuring and Displaying IGMP

IGMP Commands Used in This Section

show ip igmp configuration	page 2-7
ip igmp	page 2-9
high-priority-forward	page 2-11
auto <[ethernet] <port-list>	page 2-10
blocked <[ethernet] <port-list>	page 2-10
forward <[ethernet] <port-list>	page 2-10
querier	page 2-11
show ip igmp	Refer to the section titled “Internet Group Management Protocol (IGMP) Status” in appendix B of the <i>Management and Configuration Guide</i> for your switch.
ip igmp fastleave <port-list>	page 2-14
ip igmp forcedfastleave <port-list>	page 2-17

Viewing the Current IGMP Configuration. This command lists the IGMP configuration for all VLANs configured on the switch or for a specific VLAN.

Syntax: show ip igmp config

Displays IGMP configuration for all VLANs on the switch.

show ip igmp <vid> config

Displays IGMP configuration for a specific VLAN on the switch, including per-port data.

(For IGMP operating status, refer to the section titled “Internet Group Management Protocol (IGMP) Status” in appendix B, “Monitoring and Analyzing Switch Operation” of the Management and Configuration Guide for you switch.)

For example, suppose you have the following VLAN and IGMP configurations on the switch:

VLAN ID	VLAN Name	IGMP Enabled	Forward with High Priority	Querier
1	DEFAULT_VLAN	Yes	No	No
22	VLAN-2	Yes	Yes	Yes
33	VLAN-3	No	No	No

You could use the CLI to display this data as follows:

```

ProCurve> show ip igmp config
IGMP Service
  VLAN ID      VLAN NAME      IGMP Enabled Forward with High Priority Querier
  -----
  1            DEFAULT_VLAN  Yes          No                    No
  22           VLAN-2        Yes          Yes                   Yes
  33           VLAN-3        No           No                    Yes
  
```

Figure 2-1. Example Listing of IGMP Configuration for All VLANs in the Switch

The following version of the **show ip igmp** command includes the VLAN ID (*vid*) designation, and combines the above data with the IGMP per-port configuration:

```

IGMP Configuration for the Selected VLAN
ProCurve(config)# show ip igmp 1 config
IGMP Service
  VLAN ID : 1
  VLAN NAME : DEFAULT_VLAN
  IGMP Enabled : Yes
  Forward with High Priority : No
  Querier Allowed : Yes

IGMP Configuration On the Individual Ports in the VLAN
  Port Type | IP Mcast
  -----+-----
  A1 100/1000T | Auto
  A2 100/1000T | Auto
  A3 100/1000T | Forward
  A4 100/1000T | Forward
  A5 100/1000T | Blocked
  A6 100/1000T | Blocked
  .
  .
  .
  
```

Figure 2-2. Example Listing of IGMP Configuration for A Specific VLAN

Enabling or Disabling IGMP on a VLAN. You can enable IGMP on a VLAN, along with the last-saved or default IGMP configuration (whichever was most recently set), or you can disable IGMP on a selected VLAN.

Syntax: [no] ip igmp

Enables IGMP on a VLAN. Note that this command must be executed in a VLAN context.

For example, here are methods to enable and disable IGMP on the default VLAN (VID = 1).

```
ProCurve(config)# vlan 1 ip igmp
```

Enables IGMP on VLAN 1.

```
ProCurve(vlan-1)# ip igmp
```

Same as above.

```
ProCurve(config)# no vlan 1 ip igmp
```

Disables IGMP on vlan 1.

Note

If you disable IGMP on a VLAN and then later re-enable IGMP on that VLAN, the switch restores the last-saved IGMP configuration for that VLAN. For more on how switch memory operates, refer to the chapter titled “Switch Memory and Configuration” in the *Management and Configuration Guide* for your switch.

You can also combine the ip igmp command with other IGMP-related commands, as described in the following sections.

Configuring Per-Port IGMP Traffic Filters.

Syntax: `vlan <vid> ip igmp [auto <port-list> | blocked <port-list> | forward <port-list>]`

*Used in the VLAN context, this command specifies how each port should handle IGMP traffic. (Default: **auto**.)*

Note: *Where a static multicast filter is configured on a port, and an IGMP filter created by this command applies to the same port, the IGMP filter overrides the static multicast filter for any inbound multicast traffic carrying the same multicast address as is configured in the static filter. (Refer to the section titled “Filter Types and Operation” in the “Port Traffic Controls” chapter of the Management and Configuration Guide for your switch.*

For example, suppose you wanted to configure IGMP as follows for VLAN 1 on the 100/1000T ports on a module in slot 1:

Ports A1-A2	auto	Filter multicast traffic. Forward IGMP traffic to hosts on these ports that belong to the multicast group for which the traffic is intended. (Also forward any multicast traffic through any of these ports that is connected to a multicast router.)
Ports A3-A4	forward	Forward all multicast traffic through this port.
Ports A5-A6	blocked	Drop all multicast traffic received from devices on these ports, and prevent any outgoing multicast traffic from moving through these ports.

Depending on the privilege level, you could use one of the following commands to configure IGMP on VLAN 1 with the above settings:

```
ProCurve(config)# vlan 1 ip igmp auto a1,a2 forward a3,a4  
blocked a5,a6
```

```
ProCurve(config)# ip igmp auto a1,a2 forward a3,a4 blocked  
a5,a6
```

The following command displays the VLAN and per-port configuration resulting from the above commands.

```
ProCurve> show igmp vlan 1 config
```

Configuring IGMP Traffic Priority.

Syntax: vlan < vid > ip igmp high-priority-forward

This command assigns “high” priority to IGMP traffic or returns a high-priority setting to “normal” priority. (The traffic will be serviced at its inbound priority.) (Default: normal.)

```
ProCurve(config)# vlan 1 ip igmp high-priority-forward
```

Configures high priority for IGMP traffic on VLAN 1.

```
ProCurve(vlan-1)# ip igmp high-priority-forward
```

Same as above command, but in the VLAN 1 context level.

```
ProCurve(vlan 1)# no ip igmp high-priority-forward
```

Returns IGMP traffic to “normal” priority.

```
ProCurve> show ip igmp config
```

Show command to display results of above high-priority commands.

Configuring the Querier Function.

Syntax: [no] vlan <vid> ip igmp querier

*This command disables or re-enables the ability for the switch to become querier if necessary. The **no** version of the command disables the querier function on the switch. The **show ip igmp config** command displays the current querier command. (Default Querier Capability: Enabled.)*

How IGMP Operates

The Internet Group Management Protocol (IGMP) is an internal protocol of the Internet Protocol (IP) suite. IP manages multicast traffic by using switches, multicast routers, and hosts that support IGMP. (In Hewlett-Packard's implementation of IGMP, a multicast router is not necessary as long as a switch is configured to support IGMP with the **querier** feature enabled.) A set of hosts, routers, and/or switches that send or receive multicast data streams to or from the same source(s) is termed a *multicast group*, and all devices in the group use the same multicast group address. The multicast group running version 2 of IGMP uses three fundamental types of messages to communicate:

- **Query:** A message sent from the querier (multicast router or switch) asking for a response from each host belonging to the multicast group. If a multicast router supporting IGMP is not present, then the switch must assume this function in order to elicit group membership information from the hosts on the network. (If you need to disable the querier feature, you can do so through the CLI, using the IGMP configuration MIB. See “Configuring the Querier Function” on page 2-11.)
- **Report (Join):** A message sent by a host to the querier to indicate that the host wants to be or is a member of a given group indicated in the report message.
- **Leave Group:** A message sent by a host to the querier to indicate that the host has ceased to be a member of a specific multicast group.

Note on IGMP version 3 support

When an IGMPv3 Join is received by the switch, it accepts the host request and begins to forward the IGMP traffic. This means that ports which have not joined the group and are not connected to routers or the IGMP Querier will not receive the group's multicast traffic.

The switch does not support the IGMPv3 “Exclude Source” or “Include Source” options in the Join Reports. Rather, the group is simply joined from all sources.

The switch does not support becoming a version 3 Querier. It will become a version 2 Querier in the absence of any other Querier on the network.

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. When the client is ready to leave the multicast group, it sends a Leave Group message to the network and ceases to be a group member. When the leave request is detected, the appropriate IGMP device will cease transmitting traffic for the designated multicast group through the port on which the leave request was received (as long as there are no other current members of that group on the affected port).

Thus, IGMP identifies members of a multicast group (within a subnet) and allows IGMP-configured hosts (and routers) to join or leave multicast groups.

IGMP Data. To display data showing active group addresses, reports, queries, querier access port, and active group address data (port, type, and access), refer to the section titled “Internet Group Management Protocol (IGMP) Status” in appendix B, “Monitoring and Analyzing Switch Operation” of the *Management and Configuration Guide* for your switch.).

Operation With or Without IP Addressing

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. See the following table.

Table 2-1. Comparison of IGMP Operation With and Without IP Addressing

IGMP Function Available With IP Addressing Configured on the VLAN	Available Without IP Addressing?	Operating Differences Without an IP Address
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

Multimedia Traffic Control with IP Multicast (IGMP)

How IGMP Operates

IGMP Function Available With IP Addressing Configured on the VLAN	Available Without IP Addressing?	Operating Differences Without an IP Address
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. (ProCurve 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 and Forced Fast-Leave IGMP (below).	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.

Automatic Fast-Leave IGMP

Fast-Leave IGMP. Depending on the switch model, Fast-Leave is enabled or disabled in the default configuration.

Switch Model or Series	Data-Driven IGMP Included?	IGMP Fast-Leave Setting	Default IGMP Behavior
Switch 6400cl Switch 6200yl Switch 5400zl Switch 5300xl Switch 4200vl Switch 3500yl Switch 3400cl Switch 2900 Switch 2500	Yes	Always Enabled	Drops unjoined multicast traffic except for always-forwarded traffic toward the Querier or multicast routers, and out of IGMP-forward ports. Selectively forwards joined multicast traffic.
Switch 2600 Switch 2600-PWR Switch 4100gl Switch 6108	No	Disabled in the Default Configuration	IGMP Fast-Leave disabled in the default configuration. Floods unjoined multicast traffic to all ports. Selectively forwards joined multicast traffic.

On switches that do not support Data-Driven IGMP, unregistered multicast groups are flooded to the VLAN rather than pruned. In this scenario, Fast-Leave IGMP can actually increase the problem of multicast flooding by removing the IGMP group filter before the Querier has recognized the IGMP

leave. The Querier will continue to transmit the multicast group during this short time, and because the group is no longer registered the switch will then flood the multicast group to all ports.

On ProCurve switches that do support Data-Driven IGMP (“Smart” IGMP), when unregistered multicasts are received the switch automatically filters (drops) them. Thus, the sooner the IGMP Leave is processed, the sooner this multicast traffic stops flowing.

Because of the multicast flooding problem mentioned above, the IGMP Fast-Leave feature is disabled by default on all ProCurve switches that do not support Data-Driven IGMP. (See the table above.) The feature can be enabled on these switches via an SNMP set of this object:

```
hpSwitchIgmpportForceLeaveState.<vid>.<port number>
```

However, this is not recommended as this will increase the amount of multicast flooding during the period between the client’s IGMP Leave and the Querier’s processing of that Leave. For more information on this topic refer to “Forced Fast-Leave IGMP” on page page 2-17.

Automatic Fast-Leave Operation. If a switch port has the following characteristics, then the Fast-Leave operation will apply:

1. Connected to only one end node
2. The end node currently belongs to a multicast group; i.e. is an IGMP client
3. 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 “5A”, but not on the switch port for IGMP clients “7A” and 7B, Server “7C”, and printer “7D”.

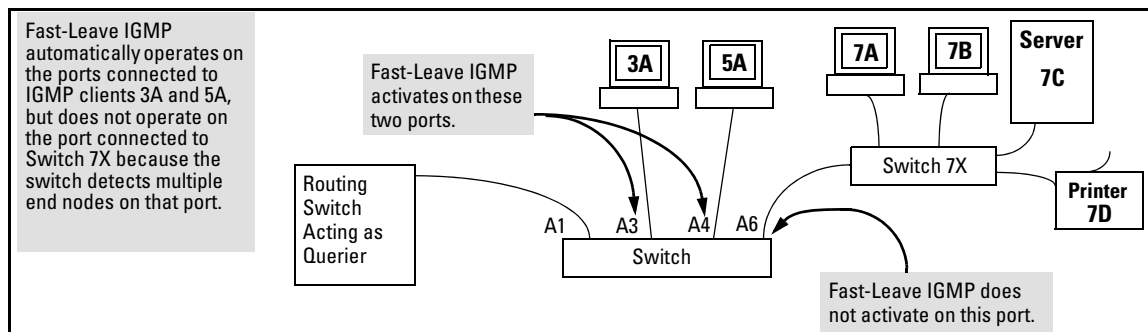


Figure 2-3. Example of Automatic Fast-Leave IGMP Criteria

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 A3, it removes the client from its IGMP table and halts multicast traffic (for that group) to port A3. 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 A3. If the switch itself is the Querier, it does not query port A3 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 A6 in figure 2-3 belong to different VLANs, Fast-Leave does not operate on port A6.

Default (Enabled) IGMP Operation Solves the “Delayed Leave” Problem. Fast-leave IGMP is enabled by default. When Fast-leave is disabled and 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 switch automatically 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 delayed leave operation means that the switch continues to transmit unnecessary multicast traffic through the port until the Querier renews multicast group status.

Configuring Fast-Leave IGMP.

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

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

Forced Fast-Leave IGMP

When enabled, 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). For example, in figure 2-3, even if you configured Forced Fast-Leave on all ports in the switch, the feature would activate only on port A6 (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.

Configuring Forced Fast-Leave IGMP

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

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

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

Forced fast-leave can be used when there are multiple devices attached to a port.

Configuring Delayed Group Flush

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

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

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

Syntax: `show igmp delayed-flush`

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

IGMP Proxy Forwarding

When a network has a border router connecting a PIM-SM domain to a PIM-DM domain, the routers that are completely within the PIM-DM domain have no way to discover multicast flows in the PIM-SM domain. When an IGMP join occurs on a router entirely within the PIM-DM domain for a flow that originates within the PIM-SM domain, it is never forwarded to the PIM-SM domain.

The IGMP proxy is a way to propagate IGMP joins across router boundaries. The proxy triggers the boundary router connected to a PIM-SM domain to query for multicast flows and forward them to the PIM-DM domain. IGMP needs to be configured on all VLAN interfaces on which the proxy is to be forwarded or received and PIM-DM must be running for the traffic to be forwarded.

You can configure an IGMP proxy on a selected VLAN that will forward IP joins (reports) and IGMP leaves to the upstream border router between the two multicast domains. You must specify the VLANs on which the proxy is enabled as well as the address of the border router to which the joins are forwarded.

How IGMP Proxy Forwarding Works

The following steps illustrate how to flood a flow from the PIM-SM domain into the PIM-DM domain when an IGMP join for that flow occurs in the PIM-DM domain (refer to figure 2-4).

1. Routing Switch 1 is configured with the IGMP proxy forwarding function to forward joins towards Border Router 1. Routing Switch 1 is also configured to forward joins from VLAN 1 toward Border Router 2, as is VLAN 4 on Routing Switch 3.
2. VLAN 2 on Routing Switch 2 is configured to forward joins toward Border Router 1.
3. When the host connected in VLAN 1 issues an IGMP join for multicast address 235.1.1.1, the join is proxied by Routing Switch 1 onto VLAN 2 and onto VLAN 4. The routing information table in Routing Switch 1 indicates that the packet to Border Router 1 and Border Router 2 is on VLAN 2 and VLAN 4, respectively.

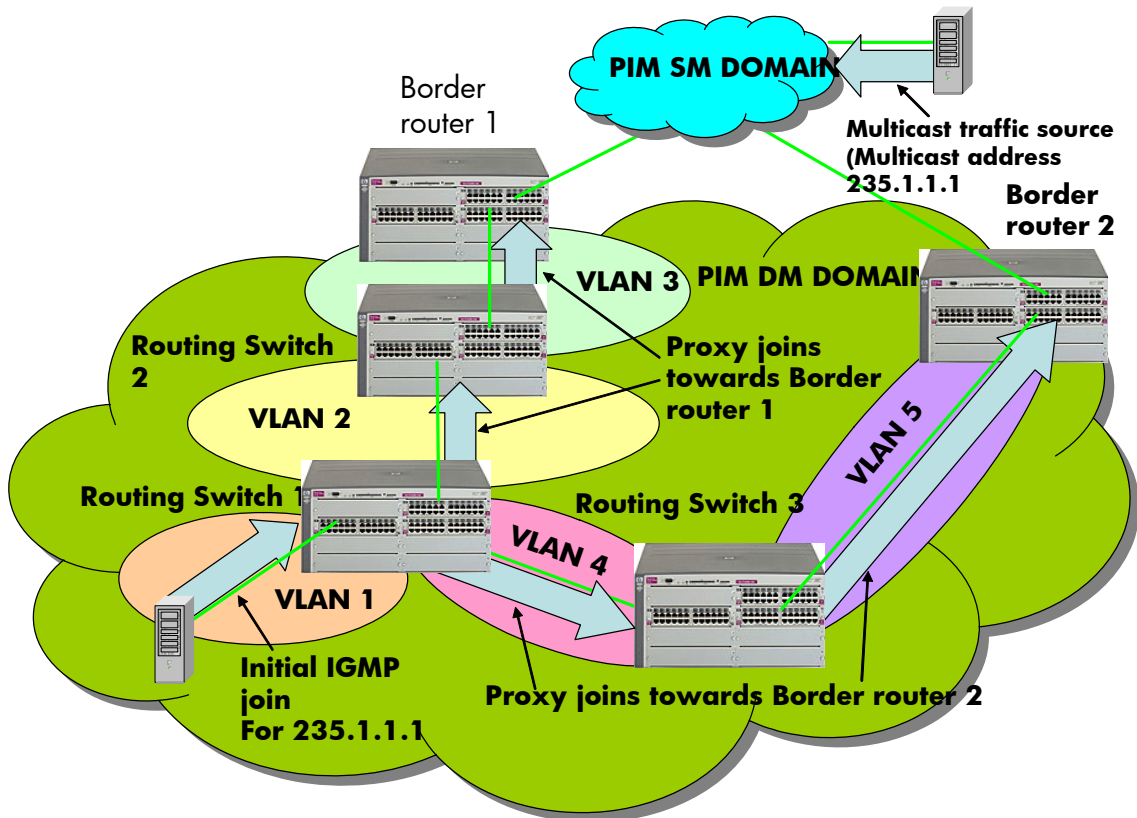


Figure 2-4. IGMP Proxy Example

4. Routing Switch 2 then proxies the IGMP join into VLAN 3, which is connected to Border Router 1.
5. Border Router 1 uses PIM-SM to find and connect to the multicast traffic for the requested traffic. The traffic is flooded into the PIM-DM network where it is routed to the original joining host.
6. Additionally, the join was proxied from Routing Switch 3 to Border Router 2. At first, both border routers will flood the traffic into the PIM-DM domain. However, PIM-DM only forwards multicasts based on the shortest reverse path back to the source of the traffic as determined by the unicast routing tables (routing FIB). Only one multicast stream is sent to the joining host. This configuration provides a redundant link in case the first link fails.

CLI Commands for IGMP Proxy Configuration

Syntax: [no] igmp-proxy-domain <domain-name> [<border-router-ip-address> <mcast-range | all>]

*Add or leave a multicast domain. The **no** form of the command is used to remove a multicast domain. All VLANs associated with the domain must first be removed for this command to work. See the **no** form of **igmp-proxy** in the VLAN context command.*

domain-name

User-defined name to associate with the PIM border router and multicast range that is being sent to toward the border router.

border-router-ip-addr

*The IP address of the border router toward which IGMP proxy packets are sent. Not required for the **no** form of the command.*

Note: The current routing FIB determines the best path towards the border router and therefore the VLAN that a proxy is sent out on.

<low-bound-ip-address | all>

The low boundary (inclusive) of the multicast address range to associate with this domain (for example, 234.0.0.1).

*If **all** is selected, the multicast addresses in the range of 224.0.0.0 - 239.255.255.255 will be included in this domain.*

Note: Addresses 224.0.0.0 - 224.0.0.255 are never used since these addresses are reserved for protocols.

<high-bound-ip-address>

The high boundary (inclusive) of the multicast address range to associate with this domain (for example, 236.1.1.1)

The following example shows the IGMP proxy border IP address (111.11.111.111) being configured.

```
ProCurve(config)# igmp-proxy-domain Bob 111.11.111.111
```

Figure 2-5. An example of the IGMP Proxy Border IP Address Command

The example below shows the lower and upper boundaries of the multicast address range associated with the domain named Bob.

```
ProCurve(config)# igmp-proxy-domain Bob 111.11.111.111 234.0.0.1  
ProCurve(config)# igmp-proxy-domain Bob 111.11.111.111 236.1.1.1
```

Figure 2-6. Setting the Lower and Upper Bounds for Multicasting

VLAN Context Command

The following command is performed when in VLAN context mode. When a query occurs on the upstream interface, an IGMP join will be sent for all multicast addresses that are currently joined on the downstream interface.

Syntax: [no] igmp-proxy <domain-name>

*Tells the VLAN which IGMP proxy domains to use with joins on the VLAN. The **no** version of the command with no domain name specified removes all domains associated with this VLAN.*

Note: Multiple different domains may be configured in the same VLAN context where the VLAN is considered the downstream interface. The domain name must exist prior to using this command to add the domain.

Note

If the unicast routing path to the specified IP address was through the VLAN specified, then no proxy IGMP would occur, that is, a proxy is not sent back out on the VLAN that the IGMP join came in on.

If no unicast route exists to the border router, then no proxy IGMP packets will be sent.

IGMP Proxy Show Command

Syntax: show igmp-proxy < entries | domains | vlans >

Shows the currently active IGMP proxy entries, domains, or vlans.

```
ProCurve(config)# show igmp-proxy entries

Total number of multicast routes: 2

Multicast Address  Border Address  VID  Multicast Domain
-----
234.43.209.12     192.168.1.1    1    George
235.22.22.12     15.43.209.1    1    SAM
226.44.3.3       192.168.1.1    2    George
```

Figure 2-7. Example Showing Active IGMP Proxy Entries

```
ProCurve(config)# show igmp-proxy domains

Total number of multicast domains: 5

Multicast Domain  Multicast Range          Border Address  Active entries
-----
George            225.1.1.1/234.43.209.12  192.168.1.1    2
SAM               235.0.0.0/239.1.1.1     15.43.209.1    1
Jane              236.234.1.1/236.235.1.1  192.160.1.2    0
Bill              ALL                       15.43.209.1    0
```

Figure 2-8. Example Showing IGMP Proxy Domains

```
ProCurve(config)# show igmp-proxy vlans

IGMP PROXY VLANs

VID          Multicast Domain  Active entries
-----
1            George           1
1            Sam              1
1            Jane           0
2            George           1
4            George           0
4            Bill            0
```

Figure 2-9. Example Showing Active IGMP Proxy VLANs

Operating Notes for IGMP Proxy Forwarding

- You can configure up to 12 multicast domains. These domains will indicate a range of multicast addresses and the IP address of the PIM-SM/PIM-DM border router.
- You must give each domain a unique name, up to 20 characters long.
- The domains may have overlapping multicast ranges.
- The IP address of the border router may be the same or different in each configured domain.
- Duplicate IGMP joins are automatically prevented, or leaves that would remove a flow currently joined by multiple hosts.
- Range overlap allows for redundant connectivity and the ability for multicasts to arrive from different border routers based on the shortest path back to the source of the traffic.
- The configured domain names must be associated with one or more VLANs for which the proxy joins are to be done.
- All routers in the path between the edge router receiving the initial IGMP packets and the border router have to be configured to forward IGMP using IGMP proxy.
- All upstream and downstream interfaces using IGMP proxy forwarding require IGMP and PIM to be enabled.

Multimedia Traffic Control with IP Multicast (IGMP)

How IGMP Operates

- You must remove all VLAN associations with the domain name before that domain name can be removed.
- The appropriate border routers must be used for each VLAN, or PIM-DM will not forward the traffic. This could occur when multiple border routers exist. It may be necessary to configure multiple overlapping domains if the multicast source address can generate the same multicast address and have different best paths to the PIM-DM domain.

Caution

Be careful to avoid configuring a IGMP forward loop, as this would leave the VLANs in a joined state forever once an initial join is sent from a host. For example, a join is issued from the host in VLAN 2 and routing switch 2 will proxy the join onto VLAN 1. Routing switch 3 will then proxy the join back onto VLAN 2 and increment its internal count of the number of joins on VLAN 2. Even after the host on VLAN 2 issues a leave, the proxy join will continue to remain and refresh itself each time a query occurs on VLAN 2. This type of loop could be created with multiple routers if an IGMP proxy is allowed to get back to the VLAN of the router that initially received the IGMP join from a host. (See figure 2-10.)

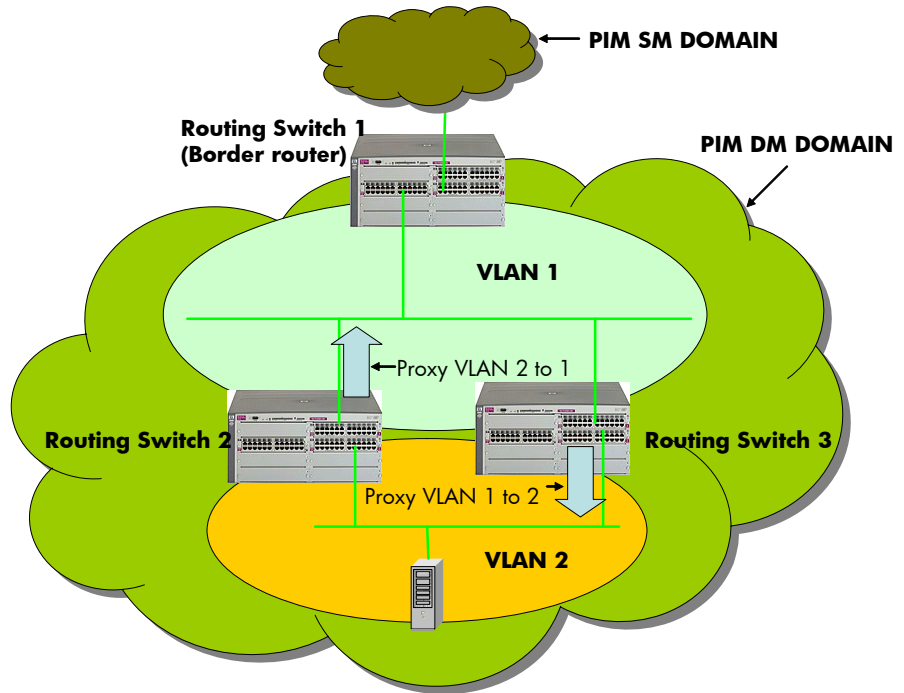


Figure 2-10. Proxy Loop Scenario

Using the Switch as Querier

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 switch's CLI to disable the Querier capability for that VLAN.

Note

A Querier is required for proper IGMP operation. For this reason, if you disable the Querier function on a 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 Querie
```

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
```

Excluding 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 pre-defined 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.

Table 2-2. 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.

Notes:

IP Multicast Filters. *This operation applies to the ProCurve 2900 switches, the Series 5400zl switches, the Series 3500yl switches, the switch 6200yl, the Series 5300xl switches, as well as the 1600M, 2400M, 2424M, 4000M, and 8000M, but not to the Series 2500, 2650, Series 4100gl, Series 4200vl, or 6108 switches (which do not have static traffic/security filters).*

IP multicast addresses occur in the range from 224.0.0.0 through 239.255.255.255 (which corresponds to the Ethernet multicast address range of 01005e-000000 through 01005e-7fffff). Where a switch has a static Traffic/Security filter configured with a “Multicast” filter type and a “Multicast Address” in this range, the switch will use the static filter unless IGMP learns of a multicast group destination in this range. In this case, IGMP dynamically takes over the filtering function for the multicast destination address(es) for as long as the IGMP group is active. If the IGMP group subsequently deactivates, the switch returns filtering control to the static filter.

Reserved Addresses Excluded from IP Multicast (IGMP) Filtering.

Traffic to IP multicast groups in the IP address range of 224.0.0.0 to 224.0.0.255 will always be flooded because addresses in this range are “well known” or “reserved” addresses. Thus, if IP Multicast is enabled and there is an IP multicast group within the reserved address range, traffic to that group will be flooded instead of filtered by the switch.

IP Routing Features

Contents

Overview of IP Routing	3-2
IP Interfaces	3-3
IP Tables and Caches	3-3
IP Global Parameters for Routing Switches	3-6
IP Interface Parameters for Routing Switches	3-8
Configuring IP Parameters for Routing Switches	3-9
Configuring IP Addresses	3-9
Configuring ARP Parameters	3-9
Configuring Forwarding Parameters	3-11
Configuring ICMP	3-13
Configuring Static IP Routes	3-15
Static Route Types	3-15
Other Sources of Routes in the Routing Table	3-16
Static IP Route Parameters	3-16
Static Route States Follow VLAN States	3-17
Configuring a Static IP Route	3-17
Displaying Static Route Information	3-19
Configuring the Default Route	3-19
Configuring IRDP	3-20
Enabling IRDP Globally	3-21
Enabling IRDP on an Individual VLAN Interface	3-21
Displaying IRDP Information	3-22
Configuring DHCP Relay	3-23
Overview	3-23
DHCP Option 82	3-23
DHCP Packet Forwarding	3-37
Minimum Requirements for DHCP Relay Operation	3-38

UDP Broadcast Forwarding	3-40
Overview	3-40
Subnet Masking for UDP Forwarding Addresses	3-41
Configuring and Enabling UDP Broadcast Forwarding	3-42
Displaying the Current IP Forward-Protocol Configuration	3-44
Operating Notes for UDP Broadcast Forwarding	3-45
Messages Related to UDP Broadcast Forwarding	3-45

Overview of IP Routing

The switches covered in this guide offer the following IP routing features, as noted:

- **IP Static Routes** – up to 16 static routes
- **IRDP** (ICMP Router Discovery Protocol) – advertises the IP addresses of the routing interfaces on this switch to directly attached host systems
- **DHCP Relay** – allows you to extend the service range of your DHCP server beyond its single local network segment

Throughout this chapter, the switches covered in this guide are referred to as “routing switches”. When IP routing is enabled on your switch, it behaves just like any other IP router.

Basic IP routing configuration consists of adding IP addresses, enabling IP routing, and enabling a route exchange protocol.

For configuring the IP addresses, refer to the chapter titled “Configuring IP Addresses” in the *Management and Configuration Guide* for your switch. The rest of this chapter describes IP routing and how to configure it in more detail. Use the information in this chapter if you need to change some of the IP parameters from their default values or you want to view configuration information or statistics.

IP Interfaces

On the routing switches, IP addresses are associated with individual VLANs. By default, there is a single VLAN (Default_VLAN) on the routing switch. In that configuration, a single IP address serves as the management access address for the entire device. If routing is enabled on the routing switch, the IP address on the single VLAN also acts as the routing interface.

Each IP address on a routing switch must be in a different sub-net. You can have only one VLAN interface that is in a given sub-net. For example, you can configure IP addresses 192.168.1.1/24 and 192.168.2.1/24 on the same routing switch, but you cannot configure 192.168.1.1/24 and 192.168.1.2/24 on the same routing switch.

You can configure multiple IP addresses on the same VLAN.

The number of IP addresses you can configure on an individual VLAN interface is 8.

You can use any of the IP addresses you configure on the routing switch for Telnet, Web management, or SNMP access, as well as for routing.

Note

All ProCurve devices support configuration and display of IP address in classical sub-net format (example: 192.168.1.1 255.255.255.0) and Classless Interdomain Routing (CIDR) format (example: 192.168.1.1/24). You can use either format when configuring IP address information. IP addresses are displayed in classical sub-net format only.

IP Tables and Caches

The following sections describe the IP tables and caches:

- ARP cache table
- IP route table
- IP forwarding cache

The software enables you to display these tables.

ARP Cache Table

The ARP cache contains entries that map IP addresses to MAC addresses. Generally, the entries are for devices that are directly attached to the routing switch.

An exception is an ARP entry for an interface-based static IP route that goes to a destination that is one or more router hops away. For this type of entry, the MAC address is either the destination device's MAC address or the MAC address of the router interface that answered an ARP request on behalf of the device, using proxy ARP.

ARP Cache. The ARP cache contains dynamic (learned) entries. The software places a dynamic entry in the ARP cache when the routing switch learns a device's MAC address from an ARP request or ARP reply from the device.

The software can learn an entry when the switch or routing switch receives an ARP request from another IP forwarding device or an ARP reply. Here is an example of a dynamic entry:

	IP Address	MAC Address	Type	Port
1	207.95.6.102	0800.5afc.ea21	Dynamic	6

Each entry contains the destination device's IP address and MAC address.

To configure other ARP parameters, see "Configuring ARP Parameters" on page 3-10.

IP Route Table

The IP route table contains routing paths to IP destinations.

Note

The default gateway, which you specify when you configure the basic IP information on the switch, is used only when routing is not enabled on the switch.

Routing Paths. The IP route table can receive the routing paths from the following sources:

- A directly-connected destination, which means there are no router hops to the destination
- A static IP route, which is a user-configured route

Administrative Distance. The IP route table contains the best path to a destination. When the software receives paths from more than one of the sources listed above, the software compares the administrative distance of each path and selects the path with the lowest administrative distance. The administrative distance is a protocol-independent value from 1 – 255.

The IP route table is displayed by entering the CLI command **show ip route** from any context level in the console CLI. Here is an example of an entry in the IP route table:

Destination	Network Mask	Gateway	Type	Sub-Type	Metric
1.1.0.0	255.255.0.0	99.1.1.2	connected		1

Each IP route table entry contains the destination’s IP address and sub-net mask and the IP address of the next-hop router interface to the destination. Each entry also indicates route type. The type indicates how the IP route table received the route.

To configure a static IP route, see “Configuring a Static IP Route” on page 3-18

IP Forwarding Cache

The IP forwarding cache provides a fast-path mechanism for forwarding IP packets. The cache contains entries for IP destinations. When an ProCurve routing switch has completed processing and addressing for a packet and is ready to forward the packet, the device checks the IP forwarding cache for an entry to the packet’s destination.

- If the cache contains an entry with the destination IP address, the device uses the information in the entry to forward the packet out the ports listed in the entry. The destination IP address is the address of the packet’s final destination. The port numbers are the ports through which the destination can be reached.
- If the cache does not contain an entry, the software can create an entry in the forwarding cache.

Each entry in the IP forwarding cache has an age timer. The age interval depends on the number of entries in the table. The age timer ranges from 12 seconds (full table) to 36 seconds (empty table). Entries are only aged if they are not being utilized by traffic. If you have an entry that is always being used in hardware, it will never age. If there is no traffic, it will age in 12-36 seconds. The age timer is not configurable.

Note

You cannot add static entries to the IP forwarding cache.

IP Global Parameters for Routing Switches

The following table lists the IP global parameters and the page where you can find more information about each parameter.

Table 3-1. IP Global Parameters for Routing Switches

Parameter	Description	Default	See page
Address Resolution Protocol (ARP)	A standard IP mechanism that routers use to learn the Media Access Control (MAC) address of a device on the network. The router sends the IP address of a device in the ARP request and receives the device's MAC address in an ARP reply.	Enabled	3-10
ARP age	The amount of time the device keeps a MAC address learned through ARP in the device's ARP cache. The device resets the timer to zero each time the ARP entry is refreshed and removes the entry if the timer reaches the ARP age.	Five minutes	not configurable
Proxy ARP	An IP mechanism a router can use to answer an ARP request on behalf of a host, by replying with the router's own MAC address instead of the host's.	Disabled	3-12
Time to Live (TTL)	The maximum number of routers (hops) through which a packet can pass before being discarded. Each router decreases a packet's TTL by 1 before forwarding the packet. If decreasing the TTL causes the TTL to be 0, the router drops the packet instead of forwarding it.	64 hops	Refer to the chapter titled "Configuring IP Addressing" in the <i>Management and Configuration Guide</i> .
Directed broadcast forwarding	A directed broadcast is a packet containing all ones (or in some cases, all zeros) in the host portion of the destination IP address. When a router forwards such a broadcast, it sends a copy of the packet out each of its enabled IP interfaces. Note: You also can enable or disable this parameter on an individual interface basis. See table 3-2 on page 3-9.	Disabled	3-13

IP Routing Features

Overview of IP Routing

Parameter	Description	Default	See page
ICMP Router Discovery Protocol (IRDP)	<p>An IP protocol that a router can use to advertise the IP addresses of its router interfaces to directly attached hosts. You can enable or disable the protocol at the Global CLI Config level.</p> <p>You also can enable or disable IRDP and configure the following protocol parameters on an individual VLAN interface basis at the VLAN Interface CLI Config level.</p> <ul style="list-style-type: none">• Forwarding method (broadcast or multicast)• Hold time• Maximum advertisement interval• Minimum advertisement interval• Router preference level	Disabled	3-21 3-22
Static route	An IP route you place in the IP route table.	No entries	3-16
Default network route	The router uses the default network route if the IP route table does not contain a route to the destination. Enter an explicit default route (0.0.0.0 0.0.0.0 or 0.0.0.0/0) as a static route in the IP route table.	None configured	3-20

IP Interface Parameters for Routing Switches

3-2 lists the interface-level IP parameters for routing switches.

Table 3-2. IP Interface Parameters – Routing Switches

Parameter	Description	Default	See page
IP address	A Layer 3 network interface address; separate IP addresses on individual VLAN interfaces.	None configured	*
ICMP Router Discovery Protocol (IRDP)	Locally overrides the global IRDP settings. See table 3-1 on page 3-7 for global IRDP information.	Disabled	3-22
IP helper address	The IP address of a UDP application server (such as a BootP or DHCP server) or a directed broadcast address. IP helper addresses allow the routing switch to forward requests for certain UDP applications from a client on one sub-net to a server on another subnet.	None configured	3-39

* Refer to the chapter titled “Configuring IP Addressing” in the Management and Configuration Guide for your routing switch.

Configuring IP Parameters for Routing Switches

The following sections describe how to configure IP parameters. Some parameters can be configured globally while others can be configured on individual VLAN interfaces. Some parameters can be configured globally and overridden for individual VLAN interfaces.

Note

This section describes how to configure IP parameters for routing switches. For IP configuration information when routing is not enabled, refer to the chapter titled “Configuring IP Addressing” in the *Management and Configuration Guide* for your routing switch.

Configuring IP Addresses

You can configure IP addresses on the routing switch’s VLAN interfaces. Configuring IP addresses is described in detail in the chapter titled “Configuring IP Addressing” in the *Management and Configuration Guide* for your switch.

Configuring ARP Parameters

Address Resolution Protocol (ARP) is a standard IP protocol that enables an IP routing switch to obtain the MAC address of another device’s interface when the routing switch knows the IP address of the interface. ARP is enabled by default and cannot be disabled.

How ARP Works

A routing switch needs to know a destination’s MAC address when forwarding traffic, because the routing switch encapsulates the IP packet in a Layer 2 packet (MAC layer packet) and sends the Layer 2 packet to a MAC interface on a device directly attached to the routing switch. The device can be the packet’s final destination or the next-hop router toward the destination.

The routing switch encapsulates IP packets in Layer 2 packets regardless of whether the ultimate destination is locally attached or is multiple router hops away. Since the routing switch’s IP route table and IP forwarding cache contain IP address information but not MAC address information, the routing switch cannot forward IP packets based solely on the information in the route

table or forwarding cache. The routing switch needs to know the MAC address that corresponds with the IP address of either the packet's locally attached destination or the next-hop router that leads to the destination.

For example, to forward a packet whose destination is multiple router hops away, the routing switch must send the packet to the next-hop router toward its destination, or to a default route or default network route if the IP route table does not contain a route to the packet's destination. In each case, the routing switch must encapsulate the packet and address it to the MAC address of a locally attached device, the next-hop router toward the IP packet's destination.

To obtain the MAC address required for forwarding a datagram, the routing switch does the following:

- First, the routing switch looks in the ARP cache (not the static ARP table) for an entry that lists the MAC address for the IP address. The ARP cache maps IP addresses to MAC addresses. The cache also lists the port attached to the device and, if the entry is dynamic, the age of the entry. A dynamic ARP entry enters the cache when the routing switch receives an ARP reply or receives an ARP request (which contains the sender's IP address and MAC address). A static entry enters the ARP cache from the static ARP table (which is a separate table) when the interface for the entry comes up.

To ensure the accuracy of the ARP cache, each dynamic entry has its own age timer. The timer is reset to zero each time the routing switch receives an ARP reply or ARP request containing the IP address and MAC address of the entry. If a dynamic entry reaches its maximum allowable age, the entry times out and the software removes the entry from the table. Static entries do not age out and can be removed only by you.

- If the ARP cache does not contain an entry for the destination IP address, the routing switch broadcasts an ARP request out all its IP interfaces. The ARP request contains the IP address of the destination. If the device with the IP address is directly attached to the routing switch, the device sends an ARP response containing its MAC address. The response is a unicast packet addressed directly to the routing switch. The routing switch places the information from the ARP response into the ARP cache.

ARP requests contain the IP address and MAC address of the sender, so all devices that receive the request learn the MAC address and IP address of the sender and can update their own ARP caches accordingly.

Note: The ARP request broadcast is a MAC broadcast, which means the broadcast goes only to devices that are directly attached to the routing switch. A MAC broadcast is not routed to other networks. However, some

routers, including ProCurve routing switches, can be configured to reply to ARP requests from one network on behalf of devices on another network. See “Enabling Proxy ARP” below.

Note

If the routing switch receives an ARP request packet that it is unable to deliver to the final destination because of the ARP time-out and no ARP response is received (the routing switch knows of no route to the destination address), the routing switch sends an ICMP Host Unreachable message to the source.

Enabling Proxy ARP

Proxy ARP allows a routing switch to answer ARP requests from devices on one network on behalf of devices in another network. Since ARP requests are MAC-layer broadcasts, they reach only the devices that are directly connected to the sender of the ARP request. Thus, ARP requests do not cross routers.

For example, if Proxy ARP is enabled on a routing switch connected to two sub-nets, 10.10.10.0/24 and 20.20.20.0/24, the routing switch can respond to an ARP request from 10.10.10.69 for the MAC address of the device with IP address 20.20.20.69. In standard ARP, a request from a device in the 10.10.10.0/24 sub-net cannot reach a device in the 20.20.20.0 sub-net if the sub-nets are on different network cables, and thus is not answered.

An ARP request from one sub-net can reach another sub-net when both sub-nets are on the same physical segment (Ethernet cable), since MAC-layer broadcasts reach all the devices on the segment.

Proxy ARP is disabled by default on ProCurve routing switches. To enable Proxy ARP, enter the following commands from the VLAN context level in the CLI:

```
ProCurve(config)# vlan 1
ProCurve(vlan-1)# ip proxy-arp
```

To again disable IP proxy ARP, enter the following command:

```
ProCurve(vlan-1)# no ip proxy-arp
```

Syntax: [no] ip proxy-arp

Configuring Forwarding Parameters

The following configurable parameters control the forwarding behavior of ProCurve routing switches:

- Time-To-Live (TTL) threshold
- Forwarding of directed broadcasts

All these parameters are global and thus affect all IP interfaces configured on the routing switch.

To configure these parameters, use the procedures in the following sections.

Changing the TTL Threshold

The configuration of this parameter is covered in the chapter titled, “Configuring IP Addressing” in the *Management and Configuration Guide* for your routing switch.

Enabling Forwarding of Directed Broadcasts

A directed broadcast is an IP broadcast to all devices within a single directly-attached network or subnet. A net-directed broadcast goes to all devices on a given network. A sub-net-directed broadcast goes to all devices within a given subnet.

Note

A less common type, the all-subnets broadcast, goes to all directly-attached subnets. Forwarding for this broadcast type also is supported, but most networks use IP multicasting instead of all-subnet broadcasting.

Forwarding for all types of IP directed broadcasts is disabled by default. You can enable forwarding for all types if needed. You cannot enable forwarding for specific broadcast types.

To enable forwarding of IP directed broadcasts, enter the following CLI command:

```
ProCurve (config) # ip directed-broadcast
```

Syntax: [no] ip directed-broadcast

ProCurve software makes the forwarding decision based on the routing switch's knowledge of the destination network prefix. Routers cannot determine that a message is unicast or directed broadcast apart from the destination network prefix. The decision to forward or not forward the message is by definition only possible in the last hop router.

To disable the directed broadcasts, enter the following CLI command:

```
ProCurve(config)# no ip directed-broadcast
```

Configuring ICMP

You can configure the following ICMP limits:

- **Burst-Normal** – The maximum number of ICMP replies to send per second.
- **Reply Limit** – You can enable or disable ICMP reply rate limiting.

Disabling ICMP Messages

ProCurve devices are enabled to reply to ICMP echo messages and send ICMP Destination Unreachable messages by default.

You can selectively disable the following types of Internet Control Message Protocol (ICMP) messages:

- **Echo messages** (ping messages) – The routing switch replies to IP pings from other IP devices.
- **Destination Unreachable messages** – If the routing switch receives an IP packet that it cannot deliver to its destination, the routing switch discards the packet and sends a message back to the device that sent the packet to the routing switch. The message informs the device that the destination cannot be reached by the routing switch.
- **Address Mask replies** – You can enable or disable ICMP address mask replies.

Disabling Replies to Broadcast Ping Requests

By default, ProCurve devices are enabled to respond to broadcast ICMP echo packets, which are ping requests. You can disable response to ping requests on a global basis using the following CLI method.

To disable response to broadcast ICMP echo packets (ping requests), enter the following command:

```
ProCurve(config)# no ip icmp echo broadcast-request
```

Syntax: [no] ip icmp echo broadcast-request

If you need to re-enable response to ping requests, enter the following command:

```
ProCurve(config)# ip icmp echo broadcast-request
```

Disabling ICMP Destination Unreachable Messages

By default, when a ProCurve device receives an IP packet that the device cannot deliver, the device sends an ICMP Unreachable message back to the host that sent the packet. The following types of ICMP Unreachable messages are generated:

- Administration – The packet was dropped by the ProCurve device due to a filter or ACL configured on the device.
- Fragmentation-needed – The packet has the “Don’t Fragment” bit set in the IP Flag field, but the ProCurve device cannot forward the packet without fragmenting it.
- Host – The destination network or subnet of the packet is directly connected to the ProCurve device, but the host specified in the destination IP address of the packet is not on the network.
- Network – The ProCurve device cannot reach the network specified in the destination IP address of the packet.
- Port – The destination host does not have the destination TCP or UDP port specified in the packet. In this case, the host sends the ICMP Port Unreachable message to the ProCurve device, which in turn sends the message to the host that sent the packet.
- Protocol – The TCP or UDP protocol on the destination host is not running. This message is different from the Port Unreachable message, which indicates that the protocol is running on the host but the requested protocol port is unavailable.
- Source-route-failure – The device received a source-routed packet but cannot locate the next-hop IP address indicated in the packet’s Source-Route option.

Note

Disabling an ICMP Unreachable message type does not change the ProCurve device’s ability to forward packets. Disabling ICMP Unreachable messages prevents the device from generating or forwarding the Unreachable messages.

To disable all ICMP Unreachable messages, enter the following command:

```
ProCurve(config)# no ip icmp unreachable
```

Syntax: [no] ip icmp unreachable

Disabling ICMP Redirects

You can disable ICMP redirects on the ProCurve routing switch only on a global basis, for all the routing switch interfaces. To disable ICMP redirects globally, enter the following command at the global CONFIG level of the CLI:

```
ProCurve(config)# no ip icmp redirects
```

Syntax: [no] ip icmp redirects

Configuring Static IP Routes

This feature enables you to create static routes (and null routes) by adding such routes directly to the route table. This section describes how to add static and null routes to the IP route table.

Static Route Types

You can configure the following types of static IP routes:

- **Standard** – the static route consists of a destination network address or host, a corresponding network mask, and the IP address of the next-hop IP address.
- **Null (discard)** – the Null route consists of the destination network address or host, a corresponding network mask, and either the **reject** or **blackhole** keyword. Typically, the null route is configured as a backup route for discarding traffic if the primary route is unavailable. By default, when IP routing is enabled, a route for the 127.0.0.0/8 network is created to the null interface. Traffic to this interface is rejected (dropped). This route is for all traffic to the “loopback” network, with the single exception of traffic to the host address of the switch’s loopback interface (127.0.0.1/32). Figure 3-2 on page 3-20 illustrates the default Null route entry in the switch’s routing table.

Note

On a single routing switch you can create one static route or null route to a given destination. Multiple static or null routes to the same destination are not supported.

Other Sources of Routes in the Routing Table

The IP route table can also receive routes from these other sources:

- **Directly-connected networks:** One route is created per IP interface. When you add an IP interface, the routing switch automatically creates a route for the network the interface is in.
- **Default route:** This is a specific static route that the routing switch uses if other routes to the destination are not available. See “Configuring the Default Route” on page 3-20.

Static IP Route Parameters

When you configure a static IP route, you must specify the following parameters:

- The IP address and network mask for the route’s destination network or host.
- The route’s path, which can be one of the following:
 - the IP address of a next-hop router.
 - a “null” interface. The routing switch drops traffic forwarded to the null interface.

The routing switch also applies default values for the following routing parameters:

- **The route’s metric:** In the case of static routes, this is the value the routing switch uses when comparing a static route to routes in the IP route table from other sources to the same destination. This is a fixed metric for static IP routes, and is set to “1”.
- **The route’s administrative distance (page 3-6):** In the case of static routes, this is the value the routing switch uses to compare a static route to routes from other route sources to the same destination before placing a route in the IP route table. The default administrative distance for static IP routes is 1, but can be configured to any value in the range of 1 - 255.

The fixed metric and administrative distance values ensure that the routing switch always prefers static IP routes over routes from other sources to the same destination.

Static Route States Follow VLAN States

IP static routes remain in the IP route table only so long as the IP interface to the next-hop router is up. If the next-hop interface goes down, the software removes the static route from the IP route table. If the next-hop interface comes up again, the software adds the route back to the route table.

This feature allows the routing switch to adjust to changes in network topology. The routing switch does not continue trying to use routes on unreachable paths but instead uses routes only when their paths are reachable.

For example, the following command configures a static route to 207.95.7.0 (with a network mask of 255.255.255.0), using 207.95.6.157 as the next-hop router's IP address.

```
ProCurve(config)# ip route 207.95.7.0/24 207.95.6.157
```

A static IP route specifies the route's destination address and the next-hop router's IP address or routing switch interface through which the routing switch can reach the destination. (The route is added to the routing switch's IP route table.)

In the above example, Router A knows that 207.95.6.157 is reachable through port A2, and assumes that local interfaces within that subnet are on the same port. Router A deduces that IP interface 207.95.7.188 is also on port A2. The software automatically removes a static IP route from the route table if the next-hop VLAN used by that route becomes unavailable. When the VLAN becomes available again, the software automatically re-adds the route to the route table.

Configuring a Static IP Route

This feature includes these options:

- **Static Route:** configure a static route to a specific network or host address
- **Null Route:** configure a “null” route to discard IP traffic to a specific network or host address:
 - discard traffic for the destination, with ICMP notification to sender
 - discard traffic for the destination, without ICMP notification to sender

Syntax: [no] ip route < dest-ip-addr >/< mask-bits >
< next-hop-ip-addr | reject | blackhole | vlan > [distance]

dest-ip-addr >/< mask-bits: The route destination and network mask length for the destination IP address. Alternatively, you can enter the mask itself. For example, you can enter either **10.0.0.0/24** or **10.0.0.0 255.255.255.0** for a route destination of 10.0.0.0 255.255.255.0.

next-hop-ip-addr: This IP address is the gateway for reaching the destination. The next-hop IP address is not required to be directly reachable on a local subnet. (If the next-hop IP address is not directly reachable, the route will be added to the routing table as soon as a route to this address is learned.)

reject: Specifies a null route where IP traffic for the specified destination is discarded and an ICMP error notification is returned to the sender.

blackhole: Specifies a null route where IP traffic for the specified destination is discarded and no ICMP error notification is returned to the sender.

vlan: Specifies the destination vlan.

distance: Specifies the administrative distance to associate with a static route. If not specified, this value is set to a default of 1. For more on this topic, refer to “Administrative Distance” on page 3-6. (Range: 1 - 255)

The **no** form of the command deletes the specified route for the specified destination next-hop pair.

The following example configures two static routes for traffic delivery and identifies two other null routes for which traffic should be discarded instead of forwarded.

```
ProCurve(config)# ip route 10.10.40.0/24 10.10.10.1
ProCurve(config)# ip route 10.10.50.128/27 10.10.10.1
ProCurve(config)# ip route 10.10.20.177/32 reject
ProCurve(config)# ip route 10.10.30.0/24 blackhole
```

Configures static routes to two different network destinations using the same next-hop router IP address.

Configures a null route to drop traffic for the device at 10.50.10.177 and return an ICMP notification to the sender.

Configures a null route to drop traffic for the 10.50.10.0 network without any ICMP notification to the sender.

Figure 3-1. Examples of Configuring Static Routes

Displaying Static Route Information

The **show ip route static** command displays the current static route configuration on the routing switch. Figure 3-2 shows the configuration resulting from the static routes configured in the preceding example.

```
ProCurve(config)# show ip route static
```

IP Route Entries					
Destination	Gateway	VLAN	Type	Sub-Type	Metric Dist.
10.10.20.177/32	reject		static		1 1
10.10.40.0/24	VLAN10	10	static		1 1
10.10.50.128/27	VLAN10	10	static		1 1
10.11.30.0/24	blackhole		static		1 1
127.0.0.0/8	reject		static		0 0

This reject (default null) route is included by default.
Refer to "Static Route Types" on page 3-16

Figure 3-2. Example of Displaying the Currently Configured Static Routes

Configuring the Default Route

You can also assign the default route and enter it in the routing table. The default route is used for all traffic that has a destination network not reachable through any other IP routing table entry. For example, if 208.45.228.35 is the IP address of your ISP router, all non-local traffic could be directed to the ISP by entering this command:

```
ProCurve(config)# ip route 0.0.0.0/0 208.45.228.35
```

Configuring IRDP

The ICMP Router Discovery Protocol (IRDP) is used by ProCurve routing switches to advertise the IP addresses of its router interfaces to directly attached hosts. IRDP is enabled by default. You can enable the feature on a global basis or on an individual VLAN interface basis.

When IRDP is enabled, the routing switch periodically sends Router Advertisement messages out the IP interfaces on which the feature is enabled. The messages advertise the routing switch's IP addresses to directly attached hosts who listen for the messages. In addition, hosts can be configured to query the routing switch for the information by sending Router Solicitation messages.

Some types of hosts use the Router Solicitation messages to discover their default gateway. When IRDP is enabled on the ProCurve routing switch, the routing switch responds to the Router Solicitation messages. Some clients interpret this response to mean that the routing switch is the default gateway. If another router is actually the default gateway for these clients, leave IRDP disabled on the ProCurve routing switch.

IRDP uses the following parameters. If you enable IRDP on individual VLAN interfaces, you can configure these parameters on an individual VLAN interface basis.

- **Packet type** - The routing switch can send Router Advertisement messages as IP broadcasts or as IP multicasts addressed to IP multicast group 224.0.0.1. The default packet type is IP broadcast.
- **Hold time** - Each Router Advertisement message contains a hold time value. This value specifies the maximum amount of time the host should consider an advertisement to be valid until a newer advertisement arrives. When a new advertisement arrives, the hold time is reset. The hold time is always longer than the maximum advertisement interval. Therefore, if the hold time for an advertisement expires, the host can reasonably conclude that the router interface that sent the advertisement is no longer available. The default hold time is three times the maximum message interval.
- **Maximum message interval and minimum message interval** - when IRDP is enabled, the routing switch sends the Router Advertisement messages every 450-600 seconds by default. The time within this interval that the routing switch selects is random for each message and is not affected by traffic loads or other network factors. The random interval minimizes the probability that a host will receive Router Advertisement

messages from other routers at the same time. The interval on each IRDP-enabled routing switch interface is independent of the interval on other IRDP-enabled interfaces. The default maximum message interval is 600 seconds. The default minimum message interval is 450 seconds.

- **Preference** - If a host receives multiple Router Advertisement messages from different routers, the host selects the router that send the message with the highest preference as the default gateway. The preference can be a number from -4294967296 to 4294967295. The default is 0.

Enabling IRDP Globally

To enable IRDP globally, enter the following command:

```
ProCurve(config)# ip irdp
```

This command enables IRDP on the IP interfaces on all ports. Each port uses the default values for the IRDP parameters.

Enabling IRDP on an Individual VLAN Interface

To enable IRDP on an individual VLAN interface and configure IRDP parameters, enter commands such as the following:

```
ProCurve(config)# vlan 1  
ProCurve(vlan-1)# ip irdp maxadvertinterval 400
```

This example shows how to enable IRDP on a specific interface (VLAN 1) and change the maximum advertisement interval for Router Advertisement messages to 400 seconds.

Syntax: [no] ip irdp [broadcast | multicast] [holdtime <seconds>] [maxadvertinterval <seconds>] [minadvertinterval <seconds>] [preference <number>]

- **broadcast | multicast** - This parameter specifies the packet type the routing switch uses to send the Router Advertisement.
 - **broadcast** - The routing switch sends Router Advertisements as IP broadcasts.
 - **multicast** - The routing switch sends Router Advertisements as multi-cast packets addressed to IP multicast group 224.0.0.1. This is the default.
- **holdtime <seconds>** - This parameter specifies how long a host that receives a Router Advertisement from the routing switch should consider the advertisement to be valid. When a host receives a new Router Advertisement message from the routing switch, the host resets the hold time

for the routing switch to the hold time specified in the new advertisement. If the hold time of an advertisement expires, the host discards the advertisement, concluding that the router interface that sent the advertisement is no longer available. The value must be greater than the value of the `maxadvertinterval` parameter and cannot be greater than 9000. The default is three times the value of the `maxadvertinterval` parameter.

- **maxadvertinterval** - This parameter specifies the maximum amount of time the routing switch waits between sending Router Advertisements. You can specify a value from 1 to the current value of the `holdtime` parameter. The default is 600 seconds.
- **minadvertinterval** - This parameter specifies the minimum amount of time the routing switch can wait between sending Router Advertisements. The default is three-fourths (0.75) the value of the `maxadvertinterval` parameter. If you change the `maxadvertinterval` parameter, the software automatically adjusts the `minadvertinterval` parameter to be three-fourths the new value of the `maxadvertinterval` parameter. If you want to override the automatically configured value, you can specify an interval from 1 to the current value of the `maxadvertinterval` parameter.
- **preference < number >** - This parameter specifies the IRDP preference level of this routing switch. If a host receives Router Advertisements from multiple routers, the host selects the router interface that sent the message with the highest preference as the host's default gateway. The valid range is -4294967296 to 4294967295. The default is 0.

Displaying IRDP Information

To display IRDP information, enter `show ip irdp` from any CLI level.

```
ProCurve# show ip irdp
Status and Counters - ICMP Router Discovery Protocol

Global Status : Disabled

VLAN Name      Status   Advertising   Min int   Max int   Holdtime   Preference
-----
Address        (sec)    (sec)        (sec)
-----
DEFAULT_VLAN   Enabled  multicast     450      600      1800      0
VLAN20         Enabled  multicast     450      600      1800      0
VLAN30         Enabled  multicast     450      600      1800      0
```

Figure 3-3. Example of Output for Show IP IRDP

Configuring DHCP Relay

Overview

The Dynamic Host Configuration Protocol (DHCP) is used for configuring hosts with IP address and other configuration parameters without human intervention. The protocol is composed of three components: the DHCP client, the DHCP server, and the DHCP relay agent. The DHCP client sends broadcast request packets to the network, the DHCP servers respond with broadcast packets that offer IP parameters, such as an IP address for the client. After the client chooses the IP parameters, communication between the client and server is by unicast packets.

The function of the DHCP relay agent is to forward the DHCP messages to other subnets so that the DHCP server doesn't have to be on the same subnet as the DHCP clients. The DHCP relay agent transfers the DHCP messages from DHCP clients located on a subnet without DHCP server, to other subnets. It also relays answers from DHCP servers to DHCP clients.

DHCP Option 82

Introduction

Option 82 is called the Relay Agent Information option and is inserted by the DHCP relay agent when forwarding client-originated DHCP packets to a DHCP server. Servers recognizing the Relay Agent Information option may use the information to implement IP address or other parameter assignment policies. The DHCP Server echoes the option back verbatim to the relay agent in server-to-client replies, and the relay agent strips the option before forwarding the reply to the client.

The "Relay Agent Information" option is organized as a single DHCP option that contains one or more "sub-options" that convey information known by the relay agent. The initial sub-options are defined for a relay agent that is co-located in a public circuit access unit. These include a "circuit ID" for the incoming circuit, and a "remote ID" which provides a trusted identifier for the remote high-speed modem.

The routing switch can operate as a DHCP relay agent to enable communication between a client and a DHCP server on a different subnet. Without Option 82, DHCP operation modifies client IP address request packets to the extent needed to forward the packets to a DHCP server. Option 82 enhances this

operation by enabling the routing switch to append an *Option 82 field* to such client requests. This field includes two suboptions for identifying the routing switch (by MAC address or IP address) and the routing switch port the client is using to access the network. A DHCP server with Option 82 capability can read the appended field and use this data as criteria for selecting the IP addressing it will return to the client through the usual DHCP server response packet. This operation provides several advantages over DHCP without Option 82:

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

Note

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

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

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

Option 82 Server Support

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

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

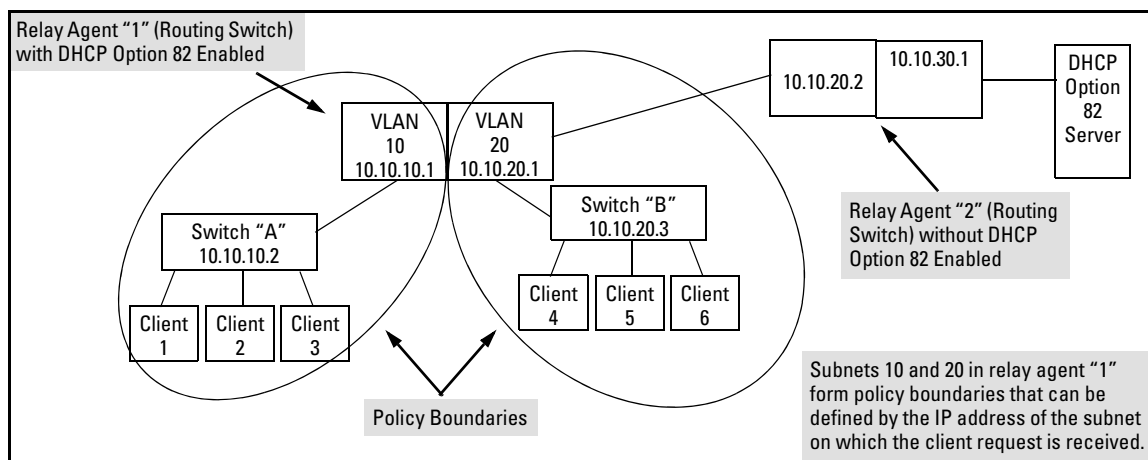


Figure 3-4. Example of a DHCP Option 82 Application

Terminology

Circuit ID: In Option 82 applications, the number of the port through which the routing switch receives a DHCP client request. On ProCurve fixed-port switches, the Circuit ID of a given port corresponds to the port number appearing on the front of the switch for that port. On ProCurve chassis switches, the port number for a given port corresponds to the internal if Index number for that port. This value is included as a suboption in an Option 82 field that the relay agent appends to a Client DHCP request before forwarding the request toward a DHCP server. (For more on Circuit ID, refer to "Circuit ID" in the bulleted list on page 3-30.)

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

DHCP relay agent: See Relay Agent.

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

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

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

Remote ID: In Option 82 applications on ProCurve switches, either the MAC address of a relay agent, or the IP address of a VLAN or subnet configured on a relay agent. This value is included as a suboption in an Option 82 field that the relay agent appends to a Client DHCP request before forwarding the request toward a DHCP server. (For more on Remote ID, refer to “Remote ID” in the bulleted list on page 3-29.)

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

General DHCP Option 82 Requirements and Operation

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

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

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

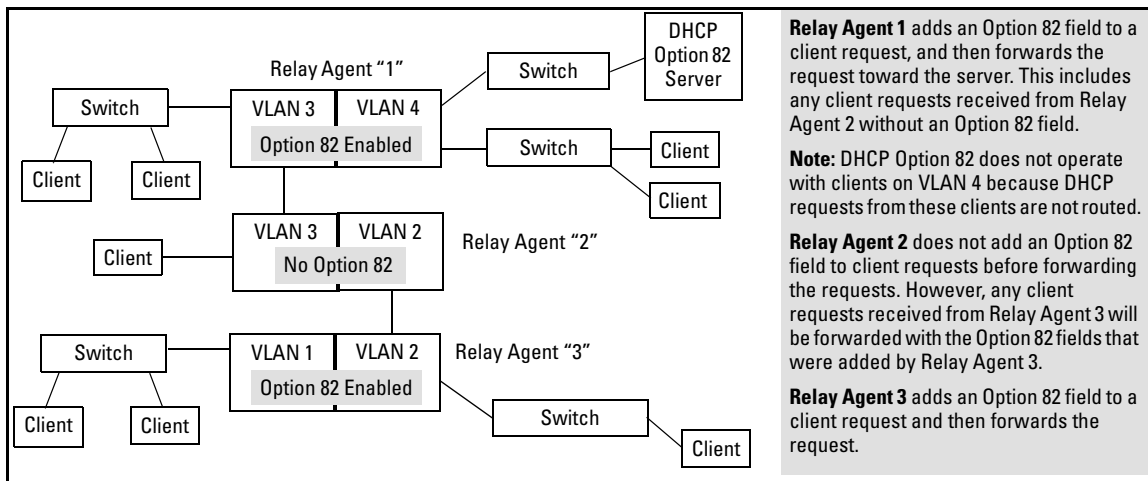


Figure 3-5. Example of DHCP Option 82 Operation in a Network with a Non-Compliant Relay Agent

Option 82 Field Content

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

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

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

```
ProCurve Switch 2900-24G(vlan-1)# show system-information

Status and Counters - General System Information

System Name       : ProCurve Switch 2900-24G
System Contact    :
System Location   :

MAC Age Time (sec) : 300

Time Zone         : 0
Daylight Time Rule : None

Software revision : T.11.XX           Base MAC Addr   : 001635-b57cc0
ROM Version       : T.11.02           Serial Number    : LP621KI005

Up Time          : 30 mins           Memory - Total  : 153,402,240
CPU Util (%)     : 23                Free            : 120,272,512

Packet - Total   : 6750
Buffers Free     : 5084
Lowest          : 5083
Missed          : 0
```




Figure 3-6. Using the CLI To View the Switch MAC Address

- **Circuit ID:** This nonconfigurable subfield identifies the port number of the physical port through which the routing switch received a given DHCP client request, and is necessary to identify if you want to configure an Option 82 DHCP server to use the Circuit ID to select a DHCP policy to assign to clients connected to the port. This number is the identity of the inbound port. On ProCurve fixed-port switches, the port number used for the Circuit ID is always the same as the physical port number shown on the front of the switch. On ProCurve chassis switches, where a dedicated, sequential block of internal port numbers are reserved for each slot, regardless of whether a slot is occupied, the circuit ID for a given port is the sequential index number for that port position in the slot. (To view the Index number assignments for ports in the routing switch, use the **walkmib ifname** command.)

```
ProCurve# walkmib ifname  
  
ifName.1 = 1  
ifName.2 = 2  
ifName.3 = 3  
ifName.4 = 4
```

Figure 3-7. Using Walkmib To Determine the Circuit ID for a Port on a ProCurve Chassis

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

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

Forwarding Policies

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

Table 3-3. Configuration Options for Managing DHCP Client Request Packets

Option 82 Configuration	DHCP Client Request Packet Inbound to the Routing Switch	
	Packet Has No Option 82 Field	Packet Includes an Option 82 Field
Append	Append an Option 82 Field	<p>Append allows the most detail in defining DHCP policy boundaries. For example, where the path from a client to the DHCP Option 82 server includes multiple relay agents with Option 82 capability, each relay agent can define a DHCP policy boundary and append its own Option 82 field to the client request packet. The server can then determine in detail the agent hops the packet took, and can be configured with a policy appropriate for any policy boundary on the path.</p> <p>Note: In networks with multiple relay agents between a client and an Option 82 server, append can be used only if the server supports multiple Option 82 fields in a client request. If the server supports only one Option 82 field in a request, consider using the keep option.</p>
Keep	Append an Option 82 Field	<p>If the relay agent receives a client request that already has one or more Option 82 fields, keep causes the relay agent to retain such fields and forward the request without adding another Option 82 field. But if the incoming client request does not already have any Option 82 fields, the relay agent appends an Option 82 field before forwarding the request. Some applications for keep include:</p> <ul style="list-style-type: none"> • The DHCP server does not support multiple Option 82 packets in a client request and there are multiple Option 82 relay agents in the path to the server. • The unusual case where DHCP clients in the network add their own Option 82 fields to their request packets and you do not want any additional fields added by relay agents. <p>This policy does not include the validate option (described in the next section) and allows forwarding of all server response packets arriving inbound on the routing switch (except those without a primary relay agent identifier.)</p>
Replace	Append an Option 82 Field	<p>Replace replaces any existing Option 82 fields from downstream relay agents (and/or the originating client) with an Option 82 field for the current relay agent. Some applications for replace include:</p> <ul style="list-style-type: none"> • The relay agent is located at a point in the network that is a DHCP policy boundary and you want to replace any Option 82 fields appended by downstream devices with an Option 82 field from the relay agent at the boundary. (This eliminates downstream Option 82 fields you do not want the server to use when determining which IP addressing policy to apply to a client request.) • In applications where the routing switch is the primary relay agent for clients that may append their own Option 82 field, you can use replace to delete these fields if you do not want them included in client requests reaching the server.
Drop	Append an Option 82 Field	<p>Drop causes the routing switch to drop an inbound client request with an Option 82 field already appended. If no Option 82 fields are present, drop causes the routing switch to add an Option 82 field and forward the request. As a general guideline, configure drop on relay agents at the edge of a network, where an inbound client request with an appended Option 82 field may be unauthorized, a security risk, or for some other reason, should not be allowed.</p>

Multiple Option 82 Relay Agents in a Client Request Path

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




Figure 3-8. Example Configured To Allow Only the Primary Relay Agent To Contribute an Option 82 Field

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




Figure 3-9. Example Configured To Allow Multiple Relay Agents To Contribute an Option 82 Field

This is an enhancement of the previous example. In this case, each hop for an accepted client request adds a new Option 82 field to the request. A DHCP server capable of using multiple Option 82 fields can be configured to use this

approach to keep a more detailed control over leased IP addresses. In this example, the primary DHCP policy boundary is at relay agent “A”, but more global policy boundaries can exist at relay agents “B” and “C”.




Figure 3-10. Example Allowing Only an Upstream Relay Agent To Contribute an Option 82 Field

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

Validation of Server Response Packets

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

Server response validation is an option you can specify when configuring Option 82 DHCP for **append**, **replace**, or **drop** operation. (Refer to “Forwarding Policies” on page 3-30.) Enabling validation on the routing switch can enhance protection against DHCP server responses that are either from untrusted sources or are carrying invalid Option 82 information.

With validation enabled, the relay agent applies stricter rules to variations in the Option 82 field(s) of incoming server responses to determine whether to forward the response to a downstream device or to drop the response due to invalid (or missing) Option 82 information. Table 3-4, below, describes relay agent management of DHCP server responses with optional validation enabled and disabled

Table 3-4. Relay Agent Management of DHCP Server Response Packets.

Response Packet Content	Option 82 Configuration	Validation Enabled on the Relay Agent	Validation Disabled (The Default)
Valid DHCP server response packet without an Option 82 field.	append, replace, or drop¹	Drop the server response packet.	Forward server response packet to a downstream device.
	keep²	Forward server response packet to a downstream device.	Forward server response packet to a downstream device.
The server response packet carries data indicating a given routing switch is the primary relay agent for the original client request, but the associated Option 82 field in the response contains a <i>Remote ID</i> and <i>Circuit ID</i> combination that did not originate with the given relay agent.	append	Drop the server response packet.	Forward server response packet to a downstream device.
	replace or drop¹	Drop the server response packet.	Drop the server response packet.
	keep²	Forward server response packet to a downstream device.	Forward server response packet to a downstream device.
The server response packet carries data indicating a given routing switch is the primary relay agent for the original client request, but the associated Option 82 field in the response contains a <i>Remote ID</i> that did not originate with the relay agent.	append	Drop the server response packet.	Forward server response packet to a downstream device.
	replace or drop¹	Drop the server response packet.	Drop the server response packet.
	keep²	Forward server response packet to a downstream device.	Forward server response packet to a downstream device.
All other server response packets ³	append, keep², replace, or drop¹	Forward server response packet to a downstream device.	Forward server response packet to a downstream device.

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

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

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

Multinetted VLANs

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

then the routing switch MAC address will be used regardless of which subnet was the source of the client request. (The MAC address is the same for all VLANs configured on the routing switch.)

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

Configuring Option 82 Operation on the Routing Switch

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

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

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

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

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

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

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

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

[validate]: *This option operates when the routing switch is configured with append, replace, or drop as a forwarding policy. With validate enabled, the routing switch applies stricter rules to an incoming Option 82 server response to determine whether to forward or drop the response. For more information, refer to "Validation of Server Response Packets" on page 3-33.*

[ip | mac]

This option specifies the remote ID suboption the routing switch will use in Option 82 fields added or appended to DHCP client packets. The choice of type depends on how you want to define DHCP policy areas in the client requests sent to the DHCP server. (Refer to “Option 82 Field Content” on page 3-28.)

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

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

Notes on Default Remote ID Selection: *Executing the Option 82 command without specifying either **ip** or **mac** configures the remote ID as the MAC address of the switch on which the packet was received from the client. The command options for viewing the routing switch MAC address are listed at the end of the “Remote ID” description that begins on page 3-28.*

Operating Notes

- This implementation of DHCP relay with Option 82 complies with the following RFCs:
 - RFC 2131
 - RFC 3046
- Moving a client to a different port allows the client to continue operating as long as the port is a member of the same VLAN as the port through which the client received its IP address. However, rebooting the client after it moves to a different port can alter the IP addressing policy the client receives if the DHCP server is configured to provide different policies to clients accessing the network through different ports.
- The IP address of the primary DHCP relay agent receiving a client request packet is automatically added to the packet, and is identified as the *giaddr* (gateway interface address). (That is, the *giaddr* is the IP address of the VLAN on which the request packet was received from the client.) For more information, refer to RFC 2131 and RFC 3046.
- DHCP request packets from multiple DHCP clients on the same relay agent port will be routed to the same DHCP server(s). Note that when using 802.1X on a switch, a port's VLAN membership may be changed by a RADIUS server responding to a client authentication request. In this case the DHCP server(s) accessible from the port may change if the VLAN assigned by the RADIUS server has different DHCP helper addresses than the VLAN used by unauthenticated clients.

- Where multiple DHCP servers are assigned to a VLAN, a DHCP client request cannot be directed to a specific server. Thus, where a given VLAN is configured for multiple DHCP servers, all of these servers should be configured with the same IP addressing policy.
- Where routing switch “A” is configured to insert its MAC address as the Remote ID in the Option 82 fields appended to DHCP client requests, and upstream DHCP servers use that MAC address as a policy boundary for assigning an IP addressing policy, then replacing switch “A” makes it necessary to reconfigure the upstream DHCP server(s) to recognize the MAC address of the replacement switch. This does not apply in the case where an upstream relay agent “B” is configured with **option 82 replace**, which removes the Option 82 field originally inserted by switch “A”.
- Relay agents without Option 82 can exist in the path between Option 82 relay agents and an Option 82 server. The agents without Option 82 will forward client requests and server responses without any effect on Option 82 fields in the packets.
- If the routing switch is not able to add an Option 82 field to a client’s DHCP request due to the message size exceeding the MTU (Maximum Transmission Unit) size, then the request is forwarded to the DHCP server without Option 82 information and an error message is logged in the switch’s Event Log.

DHCP Packet Forwarding

The DHCP relay agent on the routing switch forwards DHCP client packets to all DHCP servers that are configured in the table administrated for each VLAN.

Unicast Forwarding

The packets are forwarded using unicast forwarding if the IP address of the DHCP server is a specific host address. The DHCP relay agent sets the destination IP address of the packet to the IP address of the DHCP server and forwards the message.

Broadcast Forwarding

The packets are forwarded using broadcast forwarding if the IP address of the DHCP server is a subnet address or IP broadcast address (255.255.255.255). The DHCP relay agent sets the DHCP server IP address to broadcast IP address and will be forwarded to all VLANs with configured IP interfaces (except the source VLAN).

Minimum Requirements for DHCP Relay Operation

For the DHCP Relay agent to work, the following steps must be completed:

1. DHCP Relay is enabled on the routing switch (the default setting)
2. A DHCP server is servicing the routing switch
3. IP Routing is enabled on the routing switch
4. There is a route from the DHCP server to the routing switch and back
5. An IP Helper address is configured on the routing switch, set to the IP address of the DHCP server on the VLAN connected to the DHCP Client.

Enabling DHCP Relay

The factory-default configuration enables DHCP. However, if DHCP has been disabled, you can re-enable it at the Config CLI context level by entering this command:

```
ProCurve(config)# dhcp-relay
```

To disable the DHCP Relay function, enter the command:

```
ProCurve(config)# no dhcp-relay
```

Configuring a Helper Address

At the VLAN configuration CLI context level, enter the commands to add the DHCP server's IP address to the VLANs list. For example, to configure a helper address for VLAN 1, enter these commands:

```
ProCurve(config)# vlan 1
```

```
ProCurve(vlan-1)# ip helper-address <ip-addr>
```

To remove the DHCP server helper address, enter this command:

```
ProCurve(vlan-1)# no ip helper-address <ip-addr>
```

You can configure up to 256 IP helper addresses in the switch.

Viewing the Current DHCP Relay Configuration

Determining the DHCP Relay Setting. Use **show config** (or **show running** for the running-config file) to list the current DHCP Relay setting. Note that because DHCP Relay is enabled in the default configuration, it does not appear in these listings unless it is disabled.

```
ProCurve Switch 2900-24G(config)# show config

Startup configuration:

; J9049A Configuration Editor; Created on release #T.11.XX

hostname "ProCurve Switch 2900-24G"
module 3 type J8694A
snmp-server community "public" Unrestricted
vlan 1
    name "DEFAULT_VLAN"
    untagged 1-24,A1-A4
    ip address dhcp-bootp
    exit
no dhcp-relay
```




Figure 3-11. Example of Startup-Config Listing with DHCP-Relay Disabled

Listing the Currently Configured DHCP Helper Addresses.

Syntax: show ip helper-address < vlan-id >

This command shows the currently configured IP Helper addresses, regardless of whether DHCP-Relay is enabled. For example:



Figure 3-12. Example of Listing for IP Helper Addresses

UDP Broadcast Forwarding

Overview

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

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

Note

The number of UDP broadcast forwarding entries supported is affected by the number of IP helper addresses configured to support DHCP Relay. Refer to “Operating Notes for UDP Broadcast Forwarding” on page 3-46.

A UDP forwarding entry includes the desired UDP port number, and can be either an IP unicast address or an IP subnet broadcast address for the subnet the server is in. Thus, an incoming UDP packet carrying the configured port number will be:

- Forwarded to a specific host if a unicast server address is configured for that port number.
- Broadcast on the appropriate destination subnet if a subnet address is configured for that port number.

Note that a UDP forwarding entry for a particular UDP port number is always configured in a specific VLAN and applies only to client UDP broadcast requests received inbound on that VLAN. If the VLAN includes multiple subnets, then the entry applies to client broadcasts with that port number from any subnet in the VLAN.

For example, VLAN 1 (15.75.10.1) is configured to forward inbound UDP packets as shown in table 3-5:

Table 3-5. Example of a UDP Packet-Forwarding Environment

Interface	IP Address	Subnet Mask	Forwarding Address	UDP Port	Notes
VLAN 1	15.75.10.1	255.255.255.0	15.75.11.43	1188	Unicast address for forwarding inbound UDP packets with UDP port 1188 to a specific device on VLAN 2.
			15.75.11.255	1812	Broadcast address for forwarding inbound UDP packets with UDP port 1812 to any device in the 15.75.11.0 network.
			15.75.12.255	1813	Broadcast address for forwarding inbound UDP packets with UDP port 1813 to any device in the 15.75.12.0 network.
VLAN 2	15.75.11.1	255.255.255.0	<i>None</i>	<i>N/A</i>	Destination VLAN for UDP 1188 broadcasts from clients on VLAN 1. The device identified in the unicast forwarding address configured in VLAN 1 must be on this VLAN. Also the destination VLAN for UDP 1812 from clients on VLAN 1.
VLAN 3	15.75.12.1	255.255.255.0	<i>None</i>	<i>N/A</i>	Destination VLAN for UDP 1813 broadcasts from clients on VLAN 1.

Note If an IP server or subnet entry is invalid, a switch will not try to forward UDP packets to the configured device or subnet address.

Subnet Masking for UDP Forwarding Addresses

The subnet mask for a UDP forwarding address is the same as the mask applied to the subnet on which the inbound UDP broadcast packet is received. To forward inbound UDP broadcast packets as limited broadcasts to other subnets, use the broadcast address that covers the subnet you want to reach. For example, if VLAN 1 has an IP address of 15.75.10.1/24 (15.75.10.1 255.255.255.0), then you can configure the following unicast and limited broadcast addresses for UDP packet forwarding to subnet 15.75.11.0:

Forwarding Destination Type	IP Address
UDP Unicast to a Single Device in the 15.75.11.0 Subnet	15.75.11.X
UDP Broadcast to Subnet 15.75.11.0	15.75.11.255

Configuring and Enabling UDP Broadcast Forwarding

To configure and enable UDP broadcast forwarding on the switch:

1. Enable routing.
2. Globally enable UDP broadcast forwarding.
3. On a per-VLAN basis, configure a forwarding address and UDP port type for each type of incoming UDP broadcast you want routed to other VLANs.

Globally Enabling UDP Broadcast Forwarding

Syntax [no] ip udp-bcast-forward

*Enables or disables UDP broadcast forwarding on the router. Routing must be enabled before executing this command. Using the **no** form of this command disables any **ip forward protocol udp** commands configured in VLANs on the switch. (Default: Disabled)*

Configuring UDP Broadcast Forwarding on Individual VLANs

This command routes an inbound UDP broadcast packet received from a client on the VLAN to the unicast or broadcast address configured for the UDP port type.

Syntax [no] ip forward-protocol udp < ip-address > < port-number | port-name >

*Used in a VLAN context to configure or remove a server or broadcast address and its associated UDP port number. You can configure a maximum of 16 **forward-protocol udp** assignments in a given VLAN. The switch allows a total of 256 **forward-protocol udp** assignments across all VLANs. You can configure UDP broadcast forwarding addresses regardless of whether UDP broadcast forwarding is globally enabled on the switch. However, the feature does not operate unless globally enabled.*

— Continued on the next page. —

— Continued from the preceding page. —

< ip-address >: This can be either of the following:

- The unicast address of a destination server on another subnet. For example: 15.75.10.43.
- The broadcast address of the subnet on which a destination server operates. For example, the following address directs broadcasts to All hosts in the 15.75.11.0 subnet: 15.75.11.255.

Note: The subnet mask for a forwarded UDP packet is the same as the subnet mask for the VLAN (or subnet on a multinetted VLAN) on which the UDP broadcast packet was received from a client.

< udp-port-# >: Any UDP port number corresponding to a UDP application supported on a device at the specified unicast address or in the subnet at the specified broadcast address. For more information on UDP port numbers, refer to “TCP/UDP Port Number Ranges” on page 3-46.

< port-name >: Allows use of common names for certain well-known UDP port numbers. You can type in the specific name instead of having to recall the corresponding number:

dns: Domain Name Service (53)
nntp: Network Time Protocol (123)
netbios-ns: NetBIOS Name Service (137)
netbios-dgm: NetBIOS Datagram Service (138)
radius: Remote Authentication Dial-In User Service (1812)
radius-old: Remote Authentication Dial-In User Service (1645)
snmp: Simple Network Management Protocol (161)
snmp-trap: Simple Network Management Protocol (162)
tftp: Trivial File Transfer Protocol (69)
timep: Time Protocol (37)

For example, the following command configures the router to forward UDP broadcasts from a client on VLAN 1 for a time protocol server:

```
ProCurve(config)# ip forward-protocol udp 15.75.11.155  
timep
```


Displaying the Current IP Forward-Protocol Configuration

Syntax show ip forward-protocol [vlan < vid >]

Displays the current status of UDP broadcast forwarding and lists the UDP forwarding address(es) configured on all static VLANs in the switch or on a specific VLAN.

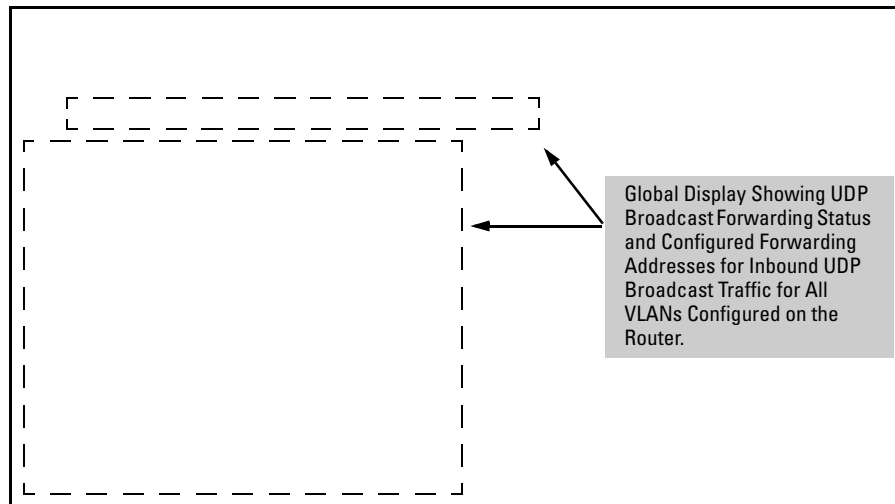


Figure 3-13. Displaying Global IP Forward-Protocol Status and Configuration




Figure 3-14. Displaying IP Forward-Protocol Status and Per-VLAN Configuration

Operating Notes for UDP Broadcast Forwarding

Maximum Number of Entries. The number of UDP broadcast entries and IP helper addresses combined can be up to 16 per VLAN, with an overall maximum of 256 on the switch. (IP helper addresses are used with the switch's DHCP Relay operation. For more information, refer to "Configuring DHCP Relay" on page 3-24.) For example, if VLAN 1 has 2 IP helper addresses configured, you can add up to 14 UDP forwarding entries in the same VLAN.

TCP/UDP Port Number Ranges. There are three ranges:

- Well-Known Ports: 0 - 1023
- Registered Ports: 1024 - 49151
- Dynamic and/or Private Ports: 49152 - 65535

For more information, including a listing of UDP/TCP port numbers, go to the *Internet Assigned Numbers Authority* (IANA) website at:

www.iana.org

Then click on:

Protocol Number Assignment Services

P (Under "Directory of General Assigned Numbers" heading)

Port Numbers

Messages Related to UDP Broadcast Forwarding

Message	Meaning
<code>udp-bcast-forward: IP Routing support must be enabled first.</code>	Appears in the CLI if an attempt to enable UDP broadcast forwarding has been made without IP routing being enabled first. Enable IP routing, then enable UDP broadcast forwarding.
<code>UDP broadcast forwarder feature enabled</code>	UDP broadcast forwarding has been globally enabled on the router. Appears in the Event Log and, if configured, in SNMP traps.
<code>UDP broadcast forwarder feature disabled</code>	UDP broadcast forwarding has been globally disabled on the router. This action does not prevent you from configuring UDP broadcast forwarding addresses, but does prevent UDP broadcast forwarding operation. Appears in the Event Log and, if configured, in SNMP traps.
<code>UDP broadcast forwarder must be disabled first.</code>	Appears in the CLI if you attempt to disable routing while UDP forwarding is enabled on the switch.

Index

A

address

IP ... 3-10

ARP

cache ... 3-5

cache table ... 3-5

configuring parameters ... 3-10

how it works ... 3-10

proxy ... 3-12

assigning

IP address ... 3-10

auto port setting ... 2-5

B

blocked port

from IGMP operation ... 2-5

broadcast traffic

enabling forwarding of directed ... 3-13

C

caches

ARP ... 3-5

IP forwarding ... 3-6

CIDR ... 3-10

configuration

ARP parameters ... 3-10

default route ... 3-20

DHCP relay ... 3-24

ICMP ... 3-14

IP routing forwarding parameters ... 3-13

IP routing parameters ... 3-10

IRDP ... 3-21

static IP routes ... 3-16, 3-18

D

default route ... 3-20

Depending ... 2-14

DHCP relay

configuration ... 3-24

enabling ... 3-39

helper address ... 3-39

minimum requirements ... 3-39

Option 82 ... 3-24

circuit ID ... 3-26, 3-30

compliance ... 3-25

configuring operation ... 3-36

field content ... 3-28

forwarding policies ... 3-27, 3-30

invalid field ... 3-25

multinetted VLANs ... 3-34

multiple relay agents ... 3-32

operation ... 3-28

policy boundary ... 3-26

relay agent ... 3-27

remote ID ... 3-27, 3-29

requirements ... 3-27

secondary relay agent ... 3-27

server response ... 3-25

server support ... 3-26

validating server response packets ... 3-33

primary relay agent ... 3-27

directed broadcasts ... 3-13

displaying information

IRDP ... 3-23

E

Exclude Source

See IGMP.

F

filters

effect of IGMP ... 2-28

maximum allowed ... 2-6

forwarding

directed broadcasts ... 3-13

forwarding parameters, IP routing

configuring ... 3-13

forwarding port, IGMP ... 2-5

H

helper address for DHCP relay ... 3-39

I

IANA ... 3-46

ICMP

- configuring ... 3-14
 - disabling messages ... 3-14
 - IGMP
 - benefits ... 2-3
 - configure per VLAN ... 2-5
 - effect on filters ... 2-28
 - Exclude Source ... 2-12
 - Fast Leave ... 2-14
 - high-priority forwarding ... 2-5
 - Include Source ... 2-12
 - IP multicast address range ... 2-28
 - leave group ... 2-12
 - maximum address count ... 2-6
 - multicast group ... 2-12
 - multimedia ... 2-3
 - operation ... 2-12, 2-13
 - port states ... 2-5
 - proxy
 - forward loop ... 2-24
 - forwarding ... 2-18
 - forwarding commands ... 2-20
 - show command ... 2-22
 - vlan context command ... 2-21
 - query ... 2-12
 - report ... 2-12
 - status ... 2-13
 - traffic ... 2-5
 - Version 3 ... 2-12
 - Include Source
 - See* IGMP.
 - interface
 - VLAN
 - enabling IRDP ... 3-22
 - IP address
 - assigning ... 3-10
 - CIDR notation ... 3-10
 - IP forwarding cache ... 3-6
 - IP global parameters ... 3-7
 - IP interface parameters ... 3-9
 - IP route table ... 3-5
 - IP routing
 - ARP cache table ... 3-5
 - changing ARP parameters ... 3-10
 - configuring static routes ... 3-16
 - default route ... 3-20
 - DHCP relay configuration ... 3-24
 - directed broadcasts ... 3-13
 - forwarding cache ... 3-6
 - forwarding parameters ... 3-13
 - global parameters ... 3-7
 - ICMP
 - configuration ... 3-14
 - disabling messages ... 3-14
 - interface parameters ... 3-9
 - IRDP configuration ... 3-21
 - overview ... 3-3
 - parameter configuring ... 3-10
 - Proxy ARP, enabling ... 3-12
 - routing table ... 3-5
 - static route configuration ... 3-18
 - static route types ... 3-16
 - tables and caches ... 3-4
 - VLAN interface ... 3-4
- I**
- IRDP
 - configuring ... 3-21
 - displaying information ... 3-23
 - enabling globally ... 3-22
 - enabling on VLAN interface ... 3-22
- L**
- leave group
 - See* IGMP.
- M**
- multicast group
 - See* IGMP.
 - multimedia
 - See* IGMP.
 - multinetted VLANs ... 3-34
 - multiple relay agents ... 3-32
- O**
- Option 82 ... 3-24
 - circuit ID ... 3-26, 3-30
 - configuring operation ... 3-36
 - field content ... 3-28
 - forwarding policies ... 3-27, 3-30
 - invalid field ... 3-25
 - multinetted VLANs ... 3-34
 - multiple relay agents ... 3-32
 - operation ... 3-28
 - policy boundary ... 3-26
 - primary relay agent ... 3-27

- relay agent ... 3-27
- remote ID ... 3-27, 3-29
- requirements ... 3-27
- secondary relay agent ... 3-27
- server support ... 3-26
- validating server response packets ... 3-33

overview, IP routing ... 3-3

P

parameters

- IP global ... 3-7
- IP interface ... 3-9

port

- auto, IGMP ... 2-5
- blocked, IGMP ... 2-5
- forwarding, IGMP ... 2-5
- state, IGMP control ... 2-5

priority ... 2-5

Proxy ARP, enabling ... 3-12

proxy forwarding, IGMP ... 2-18

Q

query

See IGMP.

quick start ... 1-8

R

relay agent information option ... 3-24

report

See IGMP.

router, multicast, with IGMP ... 2-12

routing

- configuring static routes ... 3-16
- default route ... 3-20
- DHCP relay configuration ... 3-24
- helper address ... 3-39
- helper address, UDP ... 3-9
- IP static routes ... 3-17, 3-18
 - administrative distance ... 3-17, 3-19
 - blackhole ... 3-16, 3-19
 - configuration ... 3-19
 - default route ... 3-8, 3-17
 - default route, configuring ... 3-20
 - discard traffic ... 3-18
 - discard, ICMP notification ... 3-18

- display ... 3-20
- maximum ... 3-3
- metric ... 3-17
- null interface ... 3-17
- null route ... 3-18
- null routes ... 3-16
- one per destination ... 3-16
- reject ... 3-19
- VLAN state ... 3-18

IRDP configuration ... 3-21

null routes ... 3-16

static route types ... 3-16

routing, UDP broadcast forward

See UDP broadcast forwarding.

S

setup screen ... 1-8

static IP routes

- configuring ... 3-16, 3-18

IP routing

- static route parameters ... 3-17

route types ... 3-16

subnet ... 2-13

T

tables

- ARP cache ... 3-5

- IP ... 3-4

- IP route ... 3-5

U

UDP broadcast forwarding

- address types ... 3-41

- application ... 3-41

- configure ... 3-43

- global enable ... 3-43

- invalid entry ... 3-42

- IP helper address, effect ... 3-41

- maximum entries ... 3-41

- port-number ranges ... 3-46

- show command ... 3-45

- subnet address ... 3-41

- subnet masking ... 3-42

- UDP/TCP port number listing ... 3-46

- unicast address ... 3-41

VLAN, subnetted ... 3-41

V

VLAN

IGMP configuration ... 2-5

VLAN interface

description ... 3-4

enabling IRDP ... 3-22

IP routing parameters ... 3-9

VLANS

multineted ... 3-34

W

warranty ... 1-ii



Technical information in this document
is subject to change without notice.

© Copyright 2006
Hewlett-Packard Development Company, L.P.
Reproduction, adaptation, or translation
without prior written permission is prohibited
except as allowed under the copyright laws.

August 2006

Manual Part Number
5991-6199